

# PBS Pro<sup>TM</sup> 5.2

External Reference  
Specification



PBS Pro<sup>TM</sup>

Release 5.2

External Reference  
Specification

# Portable Batch System™ External Reference Specification

PBS-3BE01, Release: PBS Pro™ 5.2, Updated: March 28, 2002

Edited by: James Patton Jones

Contributing authors include: Albeaus Bayucan, Robert L. Henderson, James Patton Jones, Casimir Lesiak, Bhroam Mann, Bill Nitzberg, Tom Proett.

Copyright (c) 2000-2002 Veridian Systems, Inc.

All rights reserved under International and Pan-American Copyright Conventions. All rights reserved. Reproduction of this work in whole or in part without prior written permission of Veridian Systems is prohibited.

Veridian Systems is an operating company of the Veridian Corporation. For more information about Veridian, visit the corporate website at: [www.veridian.com](http://www.veridian.com).

Trademarks: OpenPBS, “PBS Pro”, “Portable Batch System” and the PBS Juggler logo are trademarks of Veridian Systems, Inc. All other trademarks are the property of their respective owners.

For more information, redistribution, licensing, or additional copies of this publication, contact:

Veridian Systems  
PBS Products Dept.  
2672 Bayshore Parkway, Suite 810  
Mountain View, CA 94043

Phone: +1 (650) 967-4675  
FAX: +1 (650) 967-3080  
URL: [www.pbspro.com](http://www.pbspro.com)  
Email: [sales@pbspro.com](mailto:sales@pbspro.com)  
[support@pbspro.com](mailto:support@pbspro.com)

For online purchases, visit: [store.pbspro.com](http://store.pbspro.com)

## Table of Contents

|   |      |
|---|------|
| PBS License Agreement .....                             | pf1  |
| Revision History .....                                  | pf2  |
| <b>1. Introduction</b> .....                            | 1-1  |
| 1.1. <b>Purpose</b> .....                               | 1-1  |
| 1.2. <b>Glossary</b> .....                              | 1-1  |
| 1.3. <b>Document Conventions</b> .....                  | 1-4  |
| 1.4. <b>Overview of The Portable Batch System</b> ..... | 1-4  |
| 1.5. <b>PBS Features</b> .....                          | 1-5  |
| 1.5.1. Interactive Batch Jobs .....                     | 1-5  |
| 1.5.2. Job Prologue and Epilogue Scripts .....          | 1-6  |
| 1.5.3. File Stage In and Stage Out .....                | 1-6  |
| 1.6. Acknowledgments .....                              | 1-8  |
| <b>2. General Specifications</b> .....                  | 2-1  |
| 2.1. <b>Attribute Types</b> .....                       | 2-1  |
| 2.2. <b>Batch Jobs</b> .....                            | 2-1  |
| 2.2.1. Public Job Attributes .....                      | 2-2  |
| 2.2.2. Privileged Job Attributes .....                  | 2-4  |
| 2.2.3. Read-Only Job Attributes .....                   | 2-4  |
| 2.2.4. Job Private Attributes .....                     | 2-6  |
| 2.2.5. Job Internal Data Items .....                    | 2-6  |
| 2.2.6. Interactive Batch Jobs .....                     | 2-6  |
| 2.3. <b>Queues</b> .....                                | 2-7  |
| 2.3.1. Queue Public Attributes .....                    | 2-7  |
| 2.3.2. Queue Read-Only Attributes .....                 | 2-10 |
| 2.4. <b>Batch Server Attributes</b> .....               | 2-10 |
| 2.4.1. Server Sub-Objects .....                         | 2-15 |
| 2.5. <b>Node Attributes</b> .....                       | 2-15 |
| 2.5.1. Node Public Attributes .....                     | 2-15 |
| 2.5.2. Node Read-Only Attributes .....                  | 2-16 |
| 2.6. <b>Reservation Attributes</b> .....                | 2-17 |
| 2.6.1. Resources Reservation Read-Only Attributes ..... | 2-18 |
| 2.7. <b>PBS Files</b> .....                             | 2-20 |
| 2.8. <b>Job Selection / Scheduling</b> .....            | 2-22 |
| 2.9. <b>General Identifiers</b> .....                   | 2-24 |
| 2.9.1. Account String .....                             | 2-24 |
| 2.9.2. Attribute Name .....                             | 2-25 |
| 2.9.3. Destination Identifiers .....                    | 2-25 |
| 2.9.4. Default Server .....                             | 2-25 |
| 2.9.6. Job Identifiers .....                            | 2-25 |
| 2.9.8. Resource Name .....                              | 2-26 |
| 2.9.10. User Name .....                                 | 2-26 |
| <b>3. Batch Server Functions</b> .....                  | 3-1  |
| 3.1. <b>Client Service Requests</b> .....               | 3-1  |
| 3.1.1. Server Management .....                          | 3-1  |
| 3.1.2. Queue Management .....                           | 3-2  |
| 3.1.3. Job Management .....                             | 3-2  |
| 3.2. <b>Server to Server Requests</b> .....             | 3-8  |
| 3.2.1. Track Job Request .....                          | 3-8  |
| 3.2.2. Synchronize Job Starts .....                     | 3-8  |
| 3.2.3. Job Dependency .....                             | 3-9  |
| 3.3. <b>Deferred Services</b> .....                     | 3-10 |
| 3.3.1. Job Scheduling .....                             | 3-10 |

|   |             |
|---|-------------|
| 3.3.2. File Staging .....                                 | 3-11        |
| 3.3.3. Job Initiation .....                               | 3-11        |
| 3.3.4. Job Routing .....                                  | 3-12        |
| 3.3.5. Job Exit .....                                     | 3-13        |
| 3.3.6. Job Aborts .....                                   | 3-13        |
| 3.3.7. Timed Events .....                                 | 3-13        |
| 3.3.8. Event Logging .....                                | 3-13        |
| 3.3.9. Accounting .....                                   | 3-14        |
| <b>3.4. Resource Management .....</b>                     | <b>3-16</b> |
| 3.4.1. Non Reservable Resources .....                     | 3-16        |
| 3.4.2. Reservable Resources .....                         | 3-17        |
| 3.4.3. Resource Limits .....                              | 3-17        |
| 3.4.4. Resource Names .....                               | 3-18        |
| <b>4. Interface Library .....</b>                         | <b>4-1</b>  |
| <b>4.1. Interface Library Overview .....</b>              | <b>4-1</b>  |
| <b>4.2. Interface Library Routines .....</b>              | <b>4-2</b>  |
| 4.2.1. Alter Job .....                                    | 4-3         |
| 4.2.2. Connect with Server .....                          | 4-5         |
| 4.2.3. Delete Job .....                                   | 4-6         |
| 4.2.4. Disconnect from Server .....                       | 4-7         |
| 4.2.5. Get Error Message Text .....                       | 4-8         |
| 4.2.6. Hold Job .....                                     | 4-9         |
| 4.2.7. Locate Job .....                                   | 4-10        |
| 4.2.8. Manager Server .....                               | 4-11        |
| 4.2.9. Move Job .....                                     | 4-13        |
| 4.2.10. Message Job .....                                 | 4-14        |
| 4.2.11. Order Job .....                                   | 4-15        |
| 4.2.12. Rerun Job .....                                   | 4-16        |
| 4.2.13. Release Job .....                                 | 4-17        |
| 4.2.14. Run Job .....                                     | 4-18        |
| 4.2.15. Select Job .....                                  | 4-19        |
| 4.2.16. Status of Selected Jobs .....                     | 4-22        |
| 4.2.17. Signal Job .....                                  | 4-25        |
| 4.2.18. Stage In Files .....                              | 4-26        |
| 4.2.19. Status Job .....                                  | 4-27        |
| 4.2.20. Status Queue .....                                | 4-29        |
| 4.2.21. Status Server .....                               | 4-31        |
| 4.2.22. Submit Job .....                                  | 4-33        |
| 4.2.23. Terminate Server .....                            | 4-36        |
| 4.2.24. Query Resource Availability .....                 | 4-37        |
| 4.2.25. Make and Release Resource Reservations .....      | 4-39        |
| 4.2.26. Node Status .....                                 | 4-40        |
| <b>5. User Commands .....</b>                             | <b>5-1</b>  |
| <b>5.1. General Specifications of User Commands .....</b> | <b>5-1</b>  |
| 5.1.1. Error Checking .....                               | 5-1         |
| 5.1.2. Directing Requests to Correct Server .....         | 5-1         |
| 5.1.3. Operands .....                                     | 5-1         |
| <b>5.2. General User Commands .....</b>                   | <b>5-1</b>  |
| 5.2.1. Alter Job .....                                    | 5-2         |
| 5.2.2. Delete Job .....                                   | 5-9         |
| 5.2.3. Hold Job .....                                     | 5-10        |
| 5.2.4. Move Job .....                                     | 5-12        |
| 5.2.5. Message Job .....                                  | 5-13        |
| 5.2.6. Order Jobs .....                                   | 5-14        |

|  |             |
|--|-------------|
| 5.2.7. Rerun Job .....   | 5-15        |
| 5.2.8. Release Job .....   | 5-16        |
| 5.2.9. Select Jobs .....   | 5-18        |
| 5.2.10. Signal Job .....   | 5-22        |
| 5.2.11. Status Jobs, Queues, or Server .....                               | 5-24        |
| 5.2.12. Submit Job .....   | 5-29        |
| 5.2.13. Convert NQS Scripts .....  | 5-40        |
| 5.2.14. BASL Compiler .....  | 5-41        |
| 5.2.15. Graphical User Interface: xpbs .....                               | 5-82        |
| 5.2.16. Graphical User Interface: xpbsmon .....                            | 5-90        |
| <b>6. Operator Commands .....</b>  | <b>6-1</b>  |
| 6.1. Initiate Server Operation .....                                       | 6-2         |
| 6.2. Initiate Execution Mini-server (MOM) Operation .....                  | 6-5         |
| 6.3. Initiate Execution Mini-server with Globus Support<br>Operation ..... | 6-11        |
| 6.4. Disable Queue .....   | 6-15        |
| 6.5. Enable Queue .....  | 6-16        |
| 6.6. Run Job .....   | 6-17        |
| 6.7. Start Queue .....   | 6-18        |
| 6.8. Stop Queue .....  | 6-19        |
| 6.9. Terminate Server Operation .....                                      | 6-20        |
| <b>7. Administrator Commands .....</b>                                     | <b>7-1</b>  |
| 7.1. PBS Job Schedulers .....  | 7-2         |
| 7.2. Node Management .....   | 7-17        |
| 7.3. Manage Server .....   | 7-18        |
| <b>8. The PBS Job Scheduler .....</b>                                      | <b>8-1</b>  |
| 8.1. Scheduler Overview .....  | 8-1         |
| 8.2. The Scheduling Cycle .....  | 8-1         |
| 8.3. The Tcl Scheduler .....   | 8-1         |
| 8.4. The C Scheduler .....   | 8-1         |
| <b>9. Resource Monitor .....</b>   | <b>9-1</b>  |
| 9.1. Resources .....   | 9-1         |
| 9.1.1. SunOS Resources .....   | 9-2         |
| 9.1.2. Digital Unix .....  | 9-3         |
| 9.1.3. FreeBSD Resources .....   | 9-3         |
| 9.1.4. SGI Resources .....   | 9-3         |
| 9.1.5. Solaris Resources .....   | 9-3         |
| 9.1.6. Fujitsu Resources .....   | 9-3         |
| 9.1.7. Cray Unicos Resources .....   | 9-4         |
| 9.1.8. Cray Unicos MK 2 Resources .....                                    | 9-5         |
| 9.1.9. IBM SP .....  | 9-5         |
| 9.2. Libraries .....   | 9-5         |
| 9.2.1. Resource Monitor Library .....                                      | 9-6         |
| 9.2.2. Task Management Library .....                                       | 9-6         |
| 9.2.3. Reliable Packet Protocol Library .....                              | 9-9         |
| <b>10. Security .....</b>  | <b>10-1</b> |
| 10.1. Authorization .....  | 10-1        |
| 10.1.1. Server Access .....  | 10-1        |
| 10.1.2. Queue Access .....   | 10-2        |
| 10.1.3. Job Access .....   | 10-3        |
| 10.1.4. Root Jobs .....  | 10-3        |
| 10.2. Authentication .....   | 10-3        |
| <b>11. Network Protocol .....</b>  | <b>11-1</b> |
| 11.1. General DIS Data Encoding .....                                      | 11-1        |

|  |       |
|--|-------|
| 11.2. <b>DIS Encoded Batch Requests</b> .....          | 11-3  |
| 11.2.2. Batch Reply Format .....                       | 11-5  |
| 11.3. <b>Description of Data Exchange Format</b> ..... | 11-6  |
| 11.4. <b>Request Format</b> .....                      | 11-6  |
| 11.5. <b>Reply Format</b> .....                        | 11-7  |
| 11.6. <b>Sequence of Subrequests</b> .....             | 11-7  |
| 11.6.1. Open Connection .....                          | 11-7  |
| 11.6.2. Service Requests .....                         | 11-7  |
| 11.6.3. Connection Closure .....                       | 11-7  |
| 11.7. <b>Description of Service Request</b> .....      | 11-7  |
| 11.7.1. Queue Job Request .....                        | 11-7  |
| 11.7.2. Delete Job Request .....                       | 11-11 |
| 11.7.3. Hold Job Request .....                         | 11-11 |
| 11.7.4. Modify Job Request .....                       | 11-12 |
| 11.7.5. Move Job Request .....                         | 11-12 |
| 11.7.6. Message Job Request .....                      | 11-12 |
| 11.7.7. Rerun Job Request .....                        | 11-12 |
| 11.7.8. Select Jobs Request .....                      | 11-12 |
| 11.7.9. Signal Job Request .....                       | 11-12 |
| 11.7.10. Status Job Request .....                      | 11-13 |
| 11.7.11. Status Queue Request .....                    | 11-13 |
| 11.7.12. Status Server Request .....                   | 11-13 |
| 11.7.13. Server Shutdown Request .....                 | 11-13 |
| 11.7.14. Locate Job Request .....                      | 11-13 |
| 11.7.15. Track Job Request .....                       | 11-13 |
| 11.7.16. Run Job .....                                 | 11-14 |
| 11.7.17. Manage Request .....                          | 11-14 |
| 11.7.18. Register Dependent Job .....                  | 11-14 |
| 12. <b>Tools</b> .....                                 | 12-1  |
| 12.1. <b>Pbs_tclsh</b> .....                           | 12-1  |
| 12.2. <b>Pbs_wish</b> .....                            | 12-1  |
| 13. <b>Task Manger</b> .....                           | 13-1  |

## 1. Introduction

### 1.1. Purpose

This document is the Portable Batch System External Reference Specification. It describes the overall design of the Portable Batch System, the **PBS**, and details its external behaviors and interfaces. User, Operator, and Administrator commands are described. The interface library which is used by the commands and also may be used to extend the functionality of **PBS** is described, as is the application level data exchange protocol (network protocol).

#### NOTICE

However, this system is currently in development and the interfaces and functionality described are subject to change during the course of development.

Suggested reading of sections of the ERS is as follows:

#### General User

Users who just wish to make use of the common features of PBS to submit, monitor and control jobs are advised to read ERS chapters 1, 2.2, 2.7, 3.2 through 3.4, and 5.

#### Programmers

Programmers wishing to add an interface to PBS to their code should read the sections listed for the general user and chapter 4.

#### Operators

Batch system operator should read the sections listed under general user above plus chapters 6 and 7.

#### Administrators

Batch system administrators or managers are advised to read the entire ERS with the possible exception of chapters 4 and 11.

The requirements for **PBS** are given in the document Portable Batch System Requirement Specifications. The internal design of each component of **PBS** is specified in the document Portable Batch System Internal Design Specification.

### 1.2. Glossary

|               |   |
|---------------|---|
| Account       | is an arbitrary character string which may have meaning to one or more hosts in the batch system. Frequently, account is used as a grouping for charging for the use of resources.  |
| Administrator | See Manager.  |
| Attribute     | is an inherent characteristic of the parent object. Typically, this is a data item whose value affects the operation or behavior of the object and is settable by owner of the object. For example, the user may supply values for attributes of a job. |
| Batch         | or batch processing, is the capability of running <i>jobs</i> outside of the interactive login session. In this document, batch implies a more complex subsystem which provides for additional control over job scheduling and resource contention.     |
| Batch Server  | is a persistent subsystem (daemon) upon a single host which provides batch processing capability.   |



|                          |   |
|--------------------------|---|
| Batch System             | is a set of batch servers that are configured for processing. The system may consist of multiple hosts, each with multiple servers.   |
| Cluster                  | A complex made up of cluster nodes.   |
| Complex                  | <p>A collection of nodes managed by one batch system. A complex may be made up of nodes that are allocated to only one job at a time or of nodes that have many jobs executing on each at once or a combination of both. A complex may be made up of a set of workstations, multiple cpu systems, or one or more parallel systems.</p> <p>A queue complex in NQS was a set of queues within a batch server. The purpose of a complex was to provide additional control over resource usage. The advanced scheduling features of PBS eliminates the requirement for queue complexes.</p> |
| Cluster Node             | <p>A node whose allocation units (Virtual processors) are allocated specifically to one job at a time (see .I "exclusive node" ), or a few jobs (see <i>temporarily-shared nodes</i>). This type of node may also be called <i>space shared</i>. Hosts that are timeshared among many jobs are called "timeshared."</p> <p>In prior versions of PBS, the entire node was allocated to a job. In this version, one or more virtual processors (see definition in this glossary) may be declared. Each virtual processor is allocated as an independent unit.</p>                         |
| Destination              | is the location within the batch system where the job is sent for processing or executing. IN PBS, a destination may uniquely define a single queue at a single batch server or it may map into many locations.   |
| Destination Identifier   | <p>is a string which names the destination. It is in two parts and has the format</p> <p style="text-align: center;">queue@server</p> <p>where server is the name of a batch server and queue is the string identifying a queue on that server.</p>   |
| Exclusive Nodes          | An exclusive node is one whose virtual processors are used by one and only one job at a time. The processors of a set of nodes is assigned exclusively to a job for the duration of that job. This is typically done to improve the performance of message passing programs.  |
| Temporarily-shared Nodes | A <i>temporarily-shared node</i> is one in which one or more of the virtual processors are temporarily shared by multiple jobs. If several jobs request multiple temporarily-shared nodes, some nodes may be allocated commonly to both job and some may be unique to one of the jobs. When a node is allocated as a temporarily-shared node, it remains so until all jobs using it are terminated. Then the node may be next allocated again for temporarily-shared use or for exclusive use.  |
| File Staging             | is the movement of files between a specified location and the execution host. See "Stage In" and "Stage Out" .  |
| Group                    | is a collection of system users (see Users). A user must be a member of a group and may be a member of more than one. Within Unix and POSIX systems, membership in a group establishes one level of privilege. Group membership is also often used to control or limit access to system resources.  |
| Hold                     | is an artificial restriction which prevents a job from being selected for processing. There are three types of holds, User which is applied by the job owner, Other (or operator) which is applied by the batch operator or administrator, and System which is applied by the system itself or the batch system ad-   |

|          |  |
|----------|--|
|          | <p>ministrator.</p>  |
| Job      | <p>or <i>batch job</i> is the basic execution object managed by the batch subsystem. A job is a collection of related processes which is managed as a whole. A job can often be thought of as a shell script. In POSIX terms, a job is a session group. A session is a processes group the member processes cannot leave.</p>  |
| Manager  | <p>or Batch System Manager is a person authorized to use all restricted capabilities of the batch system. The manager may act upon the batch system, queues, or jobs. Also called the administrator.</p>   |
| Node     | <p>Within PBS, the term Node is used to mean a computational unit, one or more of which can be allocated to a job in a shared or exclusive manner. A node is characterized by (a) one or more virtual processors, (b) a unique memory address space, (c) and operating system image, (d) one or more IP addresses, and (e) a PBS execution server (<i>pbs_mom</i>). In prior versions of PBS, a node could only be allocated as a whole. Starting with version 2.2, pieces of a node, called virtual processors, VPs, may be defined and allocated individually. This is typically done where the node is a Shared Multiple Processor, SMP, system. References to allocating a <i>node</i> may still be found and should be taken to mean allocating a virtual processor. If a node only has one virtual processor defined, then of course, allocating that processor is synonymous with allocating the node itself.</p> <p>If VPs on a set of nodes are allocated in an <i>exclusive</i> manner to a job, no other job may use that set of processors during the execution of the job to which the nodes are assigned. This is often called <i>Space Sharing</i>. If a set of VPs on a set nodes are allocated in a <i>shared</i> manner, multiple jobs may have processes executing on the set at the same time.</p> |
| Operator | <p>or Batch Operator is a person authorized to use some but not all of the restricted capabilities of the batch system.</p>  |
| Owner    | <p>of a job is the user who submitted the job to the batch system.</p>   |
| PBS      | <p>is short for <i>Portable Batch System</i>.</p>  |
| POSIX    | <p>refers to the various standards being developed by the "Technical Committee on Operating Systems and Application Environments of the IEEE Computer Society" under standard P1003. There are a number of subcommittees under POSIX, those of interest to this project were:</p> <ul style="list-style-type: none"> <li>POSIX.1 System Application Program Interface (the system calls).</li> <li>POSIX.2 The command shell language.</li> <li>POSIX.3 Test Methods</li> <li>POSIX.10 Super Computing Profile</li> <li>POSIX.12 Protocol Independent Interfaces (one of the many network working groups)</li> <li>POSIX.14 Multiprocessor Working Group</li> <li>POSIX.15 Batch Queuing Extensions. This standard has been approved as 1003.2d and is the basis of the PBS external design.</li> </ul>  |
| Queue    | <p>is a collection of jobs (or job related tasks) within the batch queuing system. Each queue has a set of associated attributes which determine what actions are performed upon each job within the queue. Typical attributes include queue name, queue priority, resource limits, destination(s) and job count limits. Selection/scheduling of jobs is implementation defined. The use of the term "queue" does not imply the ordering is "first in, first out."</p>   |

|                   |   |
|-------------------|---|
| Rerunable         | If a batch job can be terminated and its execution restarted from the beginning without harmful side effects, then the job is said to be rerunable.   |
| Stage In          | is to move a file or files to the host prior to the batch job beginning execution.  |
| Stage Out         | is to move a file or files off of the host after the batch job completes execution.   |
| Timeshared        | In our context, to timeshare is to always allow multiple jobs to run concurrently on an execution host or node. A <i>timeshared node</i> is a node on which jobs are timeshared. Often the term <i>host</i> rather than <i>node</i> is used in conjunction with timeshared, as in <i>timeshared host</i> . If the term <i>node</i> is used without the timeshared prefix, the node is a cluster node which is allocated either exclusively or temporarily-shared. |
| User              | is a user of the compute system. Each user is identified by a unique character string, the user name; and by a unique number, the user id.  |
| User ID           | is a numeric identifier uniquely assigned to each user. Privilege to access system resources and services is typically established by the user id.  |
| Virtual Processor | An allocation unit that exists within a cluster (space-shared) node. Virtual processors, VPs, are allocated to jobs on an exclusive or temporarily shared basis. Virtual processors may or may not correspond in number to the number of real processors on the node. Fewer VPs than real processors may be defined to keep the batch load on the node down. More VPs may be defined to allow for over-subscription of the cpu resources.                         |

### 1.3. Document Conventions

The following font conventions are used throughout this document.

*New terms are introduced in italicized text.*

**Names of commands, library functions, and signals are shown in bold, serified text.**

[Error values] are shown in a sans-serif typeface and [inside brackets].

*Option Argument and operands are shown as italic.*

Attribute or data item names associated with jobs, queues, or the server are shown in a bold, sans-serif typeface.

Nonspecific values of attribute or data items are shown in a sans-serif typeface.

{Symbolic Constant} values, typically to be found in header files are shown in a sans-serif typeface and {inside braces}.

Examples of formats and type-ins are in the fixed width typewriter font.

Unresolved issues or areas which may require modification as the implementation progresses are called out by a note in a quoted paragraph, left and right indents, and are headed with the phrase "Author Note:". As these issues become resolved, this document will be updated and those sections removed.

### 1.4. Overview of The Portable Batch System

In the past, Unix systems were used in a totally interactive manner. Background jobs were just processes with their input disconnected from the terminal. However, as Unix moved on to larger and larger processors, the need to be able to schedule tasks based on available resources increased in importance. The advent of networked compute servers, smaller general systems, and workstations lead to the requirement of a networked batch capability.

The purpose of the **PBS** system is to provide additional controls over initiating or scheduling execution of batch jobs; and to allow routing of those jobs between different hosts. The batch system allows a site to define and implement policy as to what types of resources and how much of each resource can be used by different jobs. The batch system also provides a mechanism with which a user can insure a job will have access to the resources required to complete.

The batch system is made up of a number of components, the server and clients such as user commands. A server component manages a number of different objects, such as queues or jobs. Each object consists of a number of data items or attributes. The attributes are characterized as *public* attributes, *read-only public* attributes, *private* attributes, and *internal* attributes.

Public attributes have values which are supplied by or can be changed by client requests. The behavior of the object changes when the value of an attribute is changed. The values of public attributes are available upon request to clients. Read-only attributes are public attributes whose values are available as status to clients, but the clients cannot change the values. Throughout this document, public and read-only attributes will be commonly referred to simply as "attributes". *Private* attributes are those data items which are permanent. They are passed as part of the object when the object changes ownership, for example when a job moves between servers. Private items are generally not made available to client programs. *Internal* data items are not visible, are not passed with the object between servers, and depend on the server implementation. Public and private attributes will be described in this ERS. Except in special cases, internal data items will not be described.

Typical interaction between the components is based upon the client - server model, with clients making (batch) requests to servers and the servers performing work on behalf of the clients. Clients do not create or modify objects directly, but depend upon the server which manages those objects.

A batch server is a persistent process or set of processes, such as a daemon. The batch server manages batch objects such as queues and jobs. It provides batch services like creating, routing, executing, modifying, or deleting jobs for batch clients. A batch server may at times request services of other servers. During that time, the server is acting in the role of a client.

User, operator, and administrator commands are batch clients. They allow users of the batch system to request batch services via the command line. While the commands may appear to accomplish certain services, they actually request and obtain the services from a batch server by means of a batch request.

The Interface Library provided as part of **PBS** supplies the interface to a server for the supplied commands and allows the development of other application clients.

## 1.5. PBS Features

This section describes some of the special features of PBS.

### 1.5.1. Interactive Batch Jobs

With a normal batch job, the input to the job is the script supplied via the `qsub` command and the output and error streams are spooled to disk files and delivered after the job completes. PBS provides support for scheduling and running jobs that require interactive user access to the input and output of the job during run time. This access is often required to run debuggers or other programs requiring feedback as part of a job that must have scheduled access to scarce hardware.

A PBS interactive job is submitted with the `qsub` command. If the `-I` option is specified on the `qsub` command line or in a `#PBS` directive in the script file [or if `-W interactive=true` is specified on the command line or in a directive] the job is a interactive job. A job may de-

termine that it is an interactive batch job by the value of the environment variable **PBS\_ENVIRONMENT**. Instead of the normal value of **PBS\_BATCH**, it will have a value of **PBS\_INTERACTIVE**. [Sessions not created by PBS do not have the **PBS\_ENVIRONMENT** variable set.]

After submitting the job to the batch system, `qsub` remains active waiting for the job to connect over the network. When the job starts, data written to standard output and error is sent to `qsub` which displays it on its terminal (`qsub`'s standard output). The `qsub` command reads its standard input and passes the data to the job as the job's standard input. The job is connected via a pseudo tty, so job control signals and special characters are processed by the job, not the local `qsub` session. Since the job's standard input is from the terminal, via `qsub`, the script is not executed as is a normal script. However, any `#PBS` directives in the script are processed by `qsub` and sent with the job. Therefore, job attributes and resources requirements may be specified in the script as with a normal batch job.

While `qsub` is waiting for the job to start, it will recognize the interrupt signal typically generated by entering `CNTL-C` on the keyboard. `Qsub` will ask if the user wishes to exit. If the user responds with `yes` (or any string starting with a 'y') `qsub` will send a request to the server asking that the batch job be aborted. `Qsub` will also periodically check on the batch job. If `qsub` finds that the job has been deleted, `qsub` will inform the user and exit.

When the job starts execution, `qsub` will inform the user and begin to relay the input, output and error streams. Control keys are passed to the job. Thus a `CNTL-Z` will suspend the job, not the `qsub` command, and `CNTL-C` will cause a interrupt signal to the job, not to `qsub`. During this time, `qsub` will three escape commands, if the line begins with:

- `~.` `Qsub` will exit. This will end the job.
- `^^Z` (`CNTL-Z`) `Qsub` will suspend itself, the job remains running. Neither input nor output is transferred. The user may issue commands to the local shell.
- `^^Y` (`CNTL-Y`) `Qsub` will suspend the part of itself that sends input to the job. This allows the user to issue commands to the local shell while still receiving the output of the job.

Note, the last two escape lines do not work if the local shell does not support job control, e.g. is the Bourne shell `sh`.

When the job terminates, `qsub` will exit returning control to the local shell.

### 1.5.2. Job Prologue and Epilogue Scripts

PBS provides for the execution of two administrative supplied scripts with each job. The prologue script is run immediately before a job is executed, the epilogue script is run immediately after. Both scripts are run with root privilege and may used to place a "banner" on the job's output or establish part of the environment for the job such as creating temporary directories or cleaning up after the job.

### 1.5.3. File Stage In and Stage Out

PBS provides for files to be staged, or moved, before and after a job is run. The user specifies a "remote" location and name of the file and the "local" name when submitting the job. The remote name includes a host name which is typically a remote host, but may be the local execution host. If the file is local to the execution host, `/bin/cp` is used to copy the file, if the file is remote, `rcp` is used.

Staging out of files occurs as part of the post job processing. The job is shown to be in exiting, 'E', state during the staging out. Once the files are staged out to their destination, they are deleted from the execution host. If the user wishes to retain the files on the execution host, s/he should link the file to second file name using the `ln` (link) command.

Staging in files is a bit more complex. A decision must be made about when to begin to stage in files for a job. The files must be available before the job executes. The amount of time that will be required to copy the files is unknown to PBS, that being a function of file size and network speed. If file in-staging is not started until the job has been selected to run when the other required resources are available, either those resources are “wasted” while the stage in occurs, or another job is started which takes the resources away from the first job, and might prevent it from running. If the files are staged in well before the job is otherwise ready to run, the files may take up valuable disk space need by running jobs.

PBS provides two ways that file in-staging can be initiated for a job. If a run request is received for a job with a requirement for staging-in files, the staging in operation is begun and when completed, the job is run. Or, a specific stage-in request may be received for a job, see `pbs_stagein(3B)`, in which case the files are staged in but the job is not run. When the job is run, it begins execution immediately because the files are already there.

In either case, if the files could not be staged-in for any reason, the job is placed into a wait state with a “execute at” time `{PBS_STAGEFAIL_WAIT}`, 30 minutes, in the future. A mail message is sent to the job owner requesting that s/he look into the problem. The reason the job is changed into wait state is to prevent the scheduler from constantly retrying the same job which likely would keep on failing.

Figure 1–1 shows the (sub)state changes for a job involving file in staging. The scheduler may note the substate of the job and chose to perform pre-staging via the `pbs_stagein()` call. The scheduler developer should carefully chose a stage in approach based on factors such as the likely source of the files, network speed, and disk capacity.

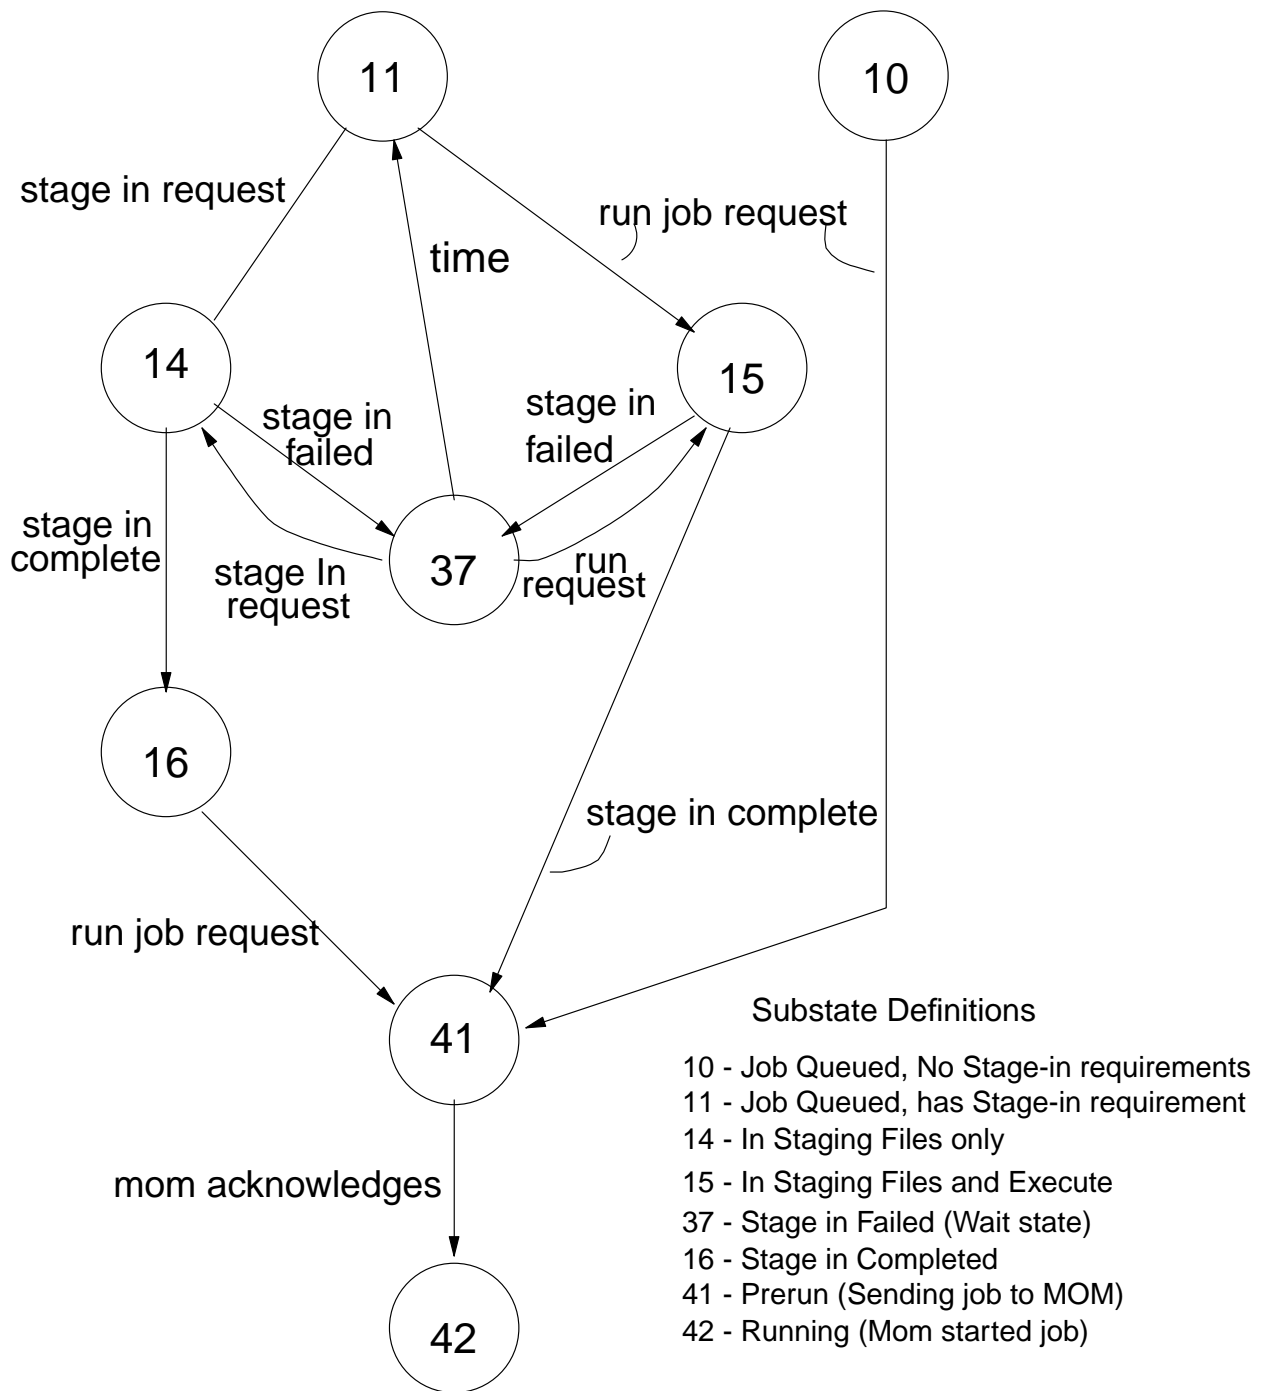


Figure 1-1: Job Substate Changes During File Stage In

1.6. Acknowledgments

Much assistance to the PBS project was given in the early days in terms of man power and ideas by both *Lawrence Livermore Nation Laboratory* and by the *National Energy Research Supercomputer Center*. Special thanks go to Bruce Kelly and Clark Streeter of NERSC, who directly assisted in the development of PBS. Additional help was provided by Kent Crispin

and Terry Heidelberg of LLNL.

The supplied code for the `mom_rcp` utility was taken from the *bsd4.4-Lite* distribution. This code is copyrighted by *The Regents of the University of California*. The complete copyright notice and right to modify and redistribute the software is contained in the source code.

The Red Hat Linux port was done by the *Pittsburgh Supercomputing Center* under funding by the *National Institute of Standards and Technology*, NIST. Special thanks go to John Kochmar and Rob Pennington of PSC for doing the port.

The ports of PBS to Digital Equipment Corporation Unix on the Alpha workstation and to HP-UX were provided by Dirk Grunwald at *University of Colorado, Boulder*.

No list of acknowledgments for PBS would possibly be complete without special recognition of the first two beta test sites and the brave individuals who were willing to try PBS. 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 tester and holds the honor of submitting more problem reports than anyone else outside of NASA. Without you two, the project would not be so successful.



[Page intentionally left blank.]

## 2. General Specifications

A server is a persistent process, a daemon, which manages several classes of objects and provides batch services. Each class of object has a set of attributes (variables) which contain information that is specific to that object. In the following sections, the objects and the services are described.

### 2.1. Attribute Types

Each attribute associated with an object has a defined data type. The following is a list of the general data types currently supported.

#### Boolean

used for true/false yes/no variables. The true state may be input as `True`, `TRUE`, `y`, `Y`, or `1`. The false state may be input as `False`, `FALSE`, `n`, `N`, or `0`. Unset boolean attributes are generally treated as if set to false.

#### Integer

used for numeric variables. The input data is a numeric string specifying a value which will fit into a long integer on the host system. Some attributes will place additional restrictions on the value range.

**Size** used for size of disk or memory related values. The input data is in the form of a numeric string with optional suffix. The suffix consist of an optional scale factor character `kKmMgGtT` and an optional byte or word indicator `bBwW`.

#### Character

used for types containing a single alphanumeric character.

#### String

used for types requiring a single null terminated character string. Additional format requirements may be placed on the string for a specific attribute.

#### Array of Strings

used where the attribute value is a series of strings. The value is input as a set of comma separated stings. The value at the human interface level (command level) often requires quoting.

**List** used for those attributes requiring a linked list of data structures. The input data is typically of the same form as the "array of strings" above.

#### Resources

is a special case of list for resource limit or resource usage items.

#### Access Control List

is a special case of list for access control lists. More information on the access control lists can be found in the ERS chapter "Security", section "Authorization".

### 2.2. Batch Jobs

A batch job is a primary object managed by a batch server. From the user's view point, a job is a file which is submitted via the **qsub** command. Typically, this file is a shell script and is interpreted by a command shell. In fact the file may contain any data and the user can request any valid program to process the file as its standard input. In addition to the shell script, a batch job consists of many attributes which affect the processing of the job. These are covered in the next section.

The server maintains an internal representation of the job and a copy of that representation as a file on disk to insure the job information is not lost across server instantiations.

Jobs are created by the server upon receipt and successful processing of a *Queue Job* batch request. Jobs are maintained by the server until (1) the job executes and terminates, (2) the job is deleted by a *Delete Job* batch request, (3) the job is moved or routed to another server, or (4) the server determines it is impossible to process the job and the job is *aborted*.

When a job is created, it is named with a *job\_identifier* by the server which created the job. The *job\_identifier* is of the form

sequence\_number.server\_name

where the *sequence\_number* is a unique number within the creating server and *server\_name* is the name of that server.

## NAME

`pbs_job_attributes` - pbs job attributes

## DESCRIPTION

### 2.2.1. Public Job Attributes

A batch job has the following public attributes shown in the following list. The attributes marked with the section symbol § are required by POSIX 1003.2d: If an attribute is unset, the indicated default value is assumed.

#### Account\_Name §

Used for accounting on some hosts. A server may not use the string, but allowances for it must be made. Format: string; default value: none, not used. [internal type: string]

#### Checkpoint §

If supported by the server implementation and the host operating system, the checkpoint attribute determines when checkpointing will be performed by PBS on behalf of the job. The legal values for checkpoint are described under the **qalter** and **qsub** commands. Format: the strings "n", "s", "c", "c=mmmm"; default value: "u", which is unspecified. [internal type: string]

**depend**The type of inter-job dependencies specified by the job owner. Format: "type:jobid[,jobid...]"; default value: no dependencies. [internal type: special, dependency]

#### Error\_Path §

The final path name for the file containing the job's standard error stream. See the **qsub** and **qalter** command description for more detail. Format: "[hostname:]path-name"; default value: (job\_name).e(job\_number). [internal type: list]

#### Execution\_Time §

The time after which the job may execute. The time is maintained in seconds since Epoch. If this time has not yet been reached, the job will not be scheduled for execution and the job is said to be in **wait** state. Format: "[CCwYY]MMDDhhmm[.ss]"; default value: time 0, no delay. [internal type: integer]

#### group\_list §

A list of *group\_names@hosts* which determines the group under which the job is run on a given host. [internal type: array of strings] When a job is to be placed into execution, the server will select a group name according to the following ordered set of rules:

1. Select the group name from the list for which the associated host name matches the name of the execution host.
2. Select the group name which has no associated host name, the "wild card name."

3. Use the login group for the user name under which the job will be run.

Format: "group\_name[@host][,group\_name[@host]...]". [internal type: array of strings]

#### Hold\_Types §

The set of holds currently applied to the job. If the set is not null, the job will not be scheduled for execution and is said to be in the **hold** state. Note, the **hold** state takes precedence over the **wait** state. Format: string made up of the letters 'u', 's', 'o'; default value: no hold. [internal type: string]

#### Job\_Name §

The name assigned to the job by the **qsub** or **qalter** command. Format: string up to 15 characters, first character must be alphabetic; default value: the base name of the job script or STDIN. [internal type: string]

#### Join\_Path §

If the Join\_Paths attribute is **TRUE**, then the job's standard error stream will be merged, inter-mixed, with the job's standard output stream and placed in the file determined by the Output\_Path attribute. The Error\_Path attribute is maintained, but ignored. Format: boolean, values accepted are "True", "TRUE", "true", "Y", "y", "1", "False", "FALSE", "false", "N", "n", "0"; default value: false. [internal type: string]

#### Keep\_Files §

If Keep\_Files contains the values "o" **KEEP\_OUTPUT** and/or "e" **KEEP\_ERROR** the corresponding streams of the batch job will be retained on the execution host upon job termination. Keep\_Files overrides the Output\_Path and Error\_Path attributes. Format: "o", "e", "oe" or "eo"; default value: no keep, return files to submission host. [internal type: string]

#### Mail\_Points §

Identifies at which state changes the server will send mail about the job. Format: string made up of the letters 'a' for abort, 'b' for beginning, and 'e' for ending; default value: 'a', send on job abort. [internal type: string]

#### Mail\_Users §

The set of users to whom mail may be sent when the job makes certain state changes. Format: "user@host[,user@host]"; default value: job owner only. [internal type: array of strings]

#### Output\_Path §

The final path name for the file containing the job's standard output stream. See the **qsub** and **qalter** command description for more detail. Format: see error\_path, default value: (job\_name).o(job\_number). [internal type: string]

#### Priority §

The job scheduling priority assigned by the user. Format: "[+|-]nnnnn"; default value: undefined. [internal type: integer]

#### Rerunable §

The rerunable flag given by the user. Format: "y" or "n", see Join\_Path; default value: y, job is rerunable. [internal type: boolean]

#### Resource\_List §

The list of resources required by the job. The resource list is a set of *name=value* strings. The meaning of *name* and *value* is server dependent. The value also establishes the limit of usage of that resource. If not set, the value for a resource may be determined by a queue or server default established by the administrator. Default value: no usage or no limit depending on specific resource. [internal type: resource]

#### Shell\_Path\_List §

A set of absolute paths of the program to process the job's script file. The list is in the format: "path[@host][,path[@host]...]". If this is null, then the user's login shell on the host of execution will be used. Default value: null, login shell. [internal type: array of

strings]

#### stagein

The list of files to be staged in prior to job execution. Format: local\_path@remote\_host:remote\_path [internal type: array of strings]

#### stageout

The list of files to be staged out after job execution. Format: local\_path@remote\_host:remote\_path [internal type: array of strings]

#### User\_List §

The list of *user@hosts* which determines the user name under which the job is run on a given host. [internal type: array of strings] When a job is to be placed into execution, the server will select a user name from the list according to the following ordered set of rules:

1. Select the user name from the list for which the associated host name matches the name of the execution host.
2. Select the user name which has no associated host name, the “wild card name.”
3. Use the Job\_Owner as the user name.

Default value: job owner name. [internal type: array of strings]

#### Variable\_List §

This is the list of environment variables passed with the *Queue Job* batch request. Format: "name=value[,name=value...]". [internal type: array of strings]

### 2.2.2. Privileged Job Attributes

The following attributes require system, manager, or operator privilege to set. They are visible to clients depending on privilege as noted.

#### comment

An attribute for displaying comments about the job from the system. Visible to any client. Format: any string; default value: none. [internal type: string]

#### sched\_hint

This attribute is present when the job is a member of a synchronous dependency set. It is set when the hold is released on the job. The value is **SYNC\_SCHED\_HINT\_FIRST** (1) when the first job of the set is released for scheduling. This is a hint that may be used by the scheduler to decrease the priority of the job. This keeps a user from attempting to “game” the scheduler. The attribute is set to **SYNC\_SCHED\_HINT\_OTHER** (2) for all other jobs in the set as they become schedulable. This should be taken as a hint by the scheduler to increase their priority to insure they will run at the same time as the earlier scheduled jobs in the set. [This attribute is viewable only by the batch administrator.] [type: integer]

### 2.2.3. Read-Only Job Attributes

The following attributes are read-only, they are established by the server and are visible to the client but cannot be set by a client. Certain ones are only visible to privileged clients (those run by the batch administrator).

**alt\_id** For a few systems, such as Irix 6.x running Array Services, the session id is insufficient to track which processes belong to the job. Where a different identifier is required, it is recorded in this attribute. If set, it will also be recorded in the end-of-job accounting record.

For Irix 6.x running Array Services, the alt\_id attribute is set to the Array Session Handle (ASH) assigned to the job. [internal type: string]

- ctime** The time that the job was created. [internal type: integer, (seconds since epoch)]
- etime** The time that the job became eligible to run, i.e. in a queued state while residing in an execution queue. [internal type: integer, (seconds since epoch)]
- exec\_host**  
If the job is running, this is set to the name of the host or hosts on which the job is executing. The format of the string is "node/N[\*C][+...]", where "node" is the name of a node, "N" is process or task slot on that node, and "C" is the number of CPUs allocated to the job. C does not appear if it is one. [internal type: string]
- egroup** If the job is queued in an execution queue, this attribute is set to the group name under which the job is to be run. [This attribute is available only to the batch administrator.] [internal type: string]
- euser** If the job is queued in an execution queue, this attribute is set to the user name under which the job is to be run. [This attribute is available only to the batch administrator.] [internal type: string]
- hashname**  
The name used as a basename for various files, such as the job file, script file, and the standard output and error of the job. [This attribute is available only to the batch administrator.] [type: string]
- interactive**  
True if the job is an interactive PBS job. Format: boolean, see Join\_Paths; default value: false. [internal type: long] Internally, the value is the port number obtained by qsub when the job was submitted.
- Job\_Owner \$**  
The login name on the submitting host of the user who submitted the batch job. [internal type: string]
- job\_state**  
The state of the job.
- E** for exiting, the job has completed execution, with or without errors, and the batch system is doing post-execution clean-up.
  - H** for Held, one or more holds have been applied to the job.
  - Q** for Queued, the job resides in a execution or routing queue pending execution or routing. It is not in **held** or **waiting** state.
  - R** for Running, the job resides in a execution queue and has been placed into execution.
  - S** for Suspend (Job running on Unicos only), the job was executing and has been suspended. The job retains its assigned resources but does not use cpu cycle or wall-time.
  - T** for Transiting, the job is in process of being routed or moved to a new destination.
  - W** for Waiting, the job is not held but the Execution\_Time attribute contains a time which has not yet been reached.
- [internal type: character]
- mtime** The time that the job was last modified, changed state, or changed locations. Internally, maintained as number of seconds since epoch. [internal type: integer]
- qtime** The time that the job entered the current queue. Internally, maintained as number of seconds since epoch. [internal type: integer]
- queue** The name of the queue in which the job currently resides. [internal type: string]
- queue\_rank**  
An ordered, non-sequential number indicating the job's position within the queue. This is provided as an aid to the scheduler. [This attribute is available to the batch manager

7  
7

only.] [internal type: integer]

queue\_type  
 An identification of the the type of queue in which the job is currently residing. This is provided as an aid to the scheduler. [This attribute is available to the batch manager only.] Format: The letter “E” or the letter “R”. [internal type: character]

resources\_used §  
 The amount of resources used by the job. This is provided as part of job status information if the job is running. [internal type: resource]

server The name of the server which is currently managing the job. [internal type: string]

session\_id  
 If the job is running, this is set to the session id of the first executing task. [internal type: integer]

substate  
 A numerical indicator of the substate of the job. The substate is used by the PBS job server internally. The attribute is visible to privileged clients, such as the scheduler. Format: interger. [internal type: long integer]  
 The values are defined in the header file job.h. See the ERS section on file staging for why it is available to the scheduler.

**SEE ALSO** 9  
 PBS ERS, qsub(1B), qalter(1B) 9

2.2.4. Job Private Attributes 9  
 The following data items are private attributes of the job. These items are a permanent part of the job object and are passed with the job between servers or between the server and the execution server, but are not passed to user clients. 9

hopcount 9  
 The hop count is maintained by the server. It is set to zero when the job is created and incremented each time the job changes destination, queue or server. The hop-count attribute is used to prevent endless routing loops and to ensure correct ordering of updates of the job’s current location for the **Job Locate** batch request. [type: integer] 9

security 9  
 Reserved for future implementation. [type: string] 9

2.2.5. Job Internal Data Items 9  
 The following data items are internal to the server representation of a job. They are specifically described here because of their importance. 9

destination 9  
 The destination\_id supplied on the **qsub** or **qmove** commands. 9

job\_substate 9  
 The secondary job state field, see “state” under Job Read-Only Attributes. As it is not visible to clients, the values are not defined in this document. 9

2.2.6. Interactive Batch Jobs 9  
 PBS supports “interactive batch jobs”. An interactive batch job is a job submitted to PBS where the standard input, output, and error streams of the job are connected to the terminal session in which qsub is run. The qsub command acts as a conduit for the communication be- 9

tween the job and the terminal session. 9

**2.3. Queues** 9

A batch queue is an object managed by a batch server. A batch queue consists of a collection of zero or more batch jobs, a set of queue attributes, private attributes, and a set of internal data items. Jobs are said to reside in the queue or be members of the queue. In spite of the name, jobs residing in a queue need not be ordered first in, first out. 9

Access to a queue is limited to the server which owns the queue. All clients gain information about a queue or jobs within a queue through batch requests to the server. 9

Two main types of queues are defined: routing queues and execution queues. The type of queue is determined by which subset of queue attributes have been assigned to it. 9

When a job resides in a routing queue, it is a candidate for routing to a new destination. Each routing queue has a list of destinations to which jobs may be routed. The new destination may be a different queue within the same server or a queue under a different server. 9

Jobs are removed from a routing queue when: 9

- The job has been successfully routed to another queue. 9
- The job has been deleted. 9
- The job has been moved to another queue. 9
- The job has been aborted by the server. 9

When a job resides in an execution queue, it is a candidate for execution. A job in execution is still a member of the execution queue from which it was selected for execution. Jobs are removed from an execution queue when: 9

- The job has executed and terminated. 9
- The job has been deleted. 9
- The job has been moved to another queue. 9
- The job has been aborted. 9

**NAME** 9

`pbs_queue_attributes` - pbs queue attributes 9

**DESCRIPTION** 9

**2.3.1. Queue Public Attributes** 9

Queue public attributes are alterable on request by a client. The client must be acting for a user with administrator (manager) or operator privilege. Certain attributes require the user to have full administrator privilege before they can be modified. The following attributes apply to both queue types: 9

`acl_group_enable` 9

Attribute which when true directs the server to use the queue group access control list `acl_groups`. Format: boolean, "TRUE", "True", "true", "Y", "y", "1", "FALSE", "False", "false", "N", "n", "0"; default value: false = disabled. [internal type: boolean] 9

`acl_groups` 9

List which allows or denies enqueueing of jobs owned by members of the listed groups. The groups in the list are groups on the server host, not submitting hosts. See section 10.1, Authorization, in the PBS External Reference Specification. Format: "[+|-]group\_name[...]"; default value: all groups allowed. [internal type: ac- 9



|  |   |
|--|---|
| cess control list]   | 9 |
| acl_host_enable  | 9 |
| Attribute which when true directs the server to use the <i>acl_hosts</i> access list. Format: boolean (see <i>acl_group_enable</i> ); default value: disabled. [internal type: boolean]  | 9 |
| acl_hosts  | 9 |
| List of hosts which may enqueue jobs in the queue. See section 10.1, Authorization, in the PBS External Reference Specification. Format: "[+ -]hostname[...]"; default value: all hosts allowed. [internal type: access control list]  | 9 |
| acl_user_enable  | 9 |
| Attribute which when true directs the server to use the the <i>acl_users</i> access list for this queue. Format: boolean (see <i>acl_group_enable</i> ); default value: disabled. [internal type: boolean]   | 9 |
| acl_users  | 9 |
| List of users allowed or denied the ability to enqueue jobs in this queue. See section 10.1, Authorization, in the PBS External Reference Specification. Format: "[+ -]user[@host][...]"; default value: all users allowed. [internal type: access control list]   | 9 |
| enabled  | 9 |
| Queue will or will not accept new jobs. When false the queue is “disabled” and will not accept jobs. Format: boolean (see <i>acl_group_enable</i> ); default value: disabled. [internal type: boolean]   | 9 |
| from_route_only  | 9 |
| When true, this queue will not accept jobs except when being routed by the server from a local routing queue. This is used to force user to submit jobs into a routing queue used to distribute jobs to other queues based on job resource limits. Requires full manager privilege to set or alter. Format: boolean; default value: disabled. [internal type: boolean] | 9 |
| max_queueable  | 9 |
| The maximum number of jobs allowed to reside in the queue at any given time. Format: integer; default value: infinite. [internal type: integer]  | 9 |
| max_running  | 9 |
| The maximum number of jobs allowed to be selected from this queue for routing or execution at any given time. For a routing queue, this is enforced, if set, by the server. For an execution queue, this attribute is advisory to the Scheduler, it is not enforced by the server. Format: integer. [internal type: integer]   | 9 |
| Priority   | 9 |
| The priority of this queue against other queues of the same type on this server. May affect job selection for execution/routing. Advisory to the Scheduler, not used by the server. Format: integer. [internal type: integer]  | 9 |
| queue_type   | 9 |
| The type of the queue: execution or route. Requires full manager privilege to set or alter. Format: "execution", "e", "route", "r". This attribute must be explicitly set. [internal type: string]   | 9 |
| resources_max  | 9 |
| The maximum amount of each resource which can be requested by a single job in this queue. The queue value superceeds any server wide maximum limit. Format: "resources_max.resource_name=value", see <i>qmgr(1B)</i> ; default value: infinite usage. [internal type: resource]  | 9 |
| resources_min  | 9 |
| The minimum amount of each resource which can be requested by a single job in  | 9 |

|  |   |
|--|---|
| this queue. Format: see resources_max, default value: zero usage. [internal type: resource]  | 9 |
| resources_default  | 9 |
| The list of default resource values which are set as limits for a job residing in this queue and for which the job did not specify a limit. Format: "resources_default.resource_name=value", see qmgr(1B); default value: none; if not set, the default limit for a job is determined by the first of the following attributes which is set: server's resources_default, queue's resources_max, server's resources_max. If none of these are set, the job will unlimited resource usage. [internal type: resource] | 9 |
| started  | 9 |
| Jobs may be scheduled for execution from this queue. When false, the queue is considered "stopped." Advisory to the Scheduler, not enforced by the server. [default value: false, but depends on scheduler interpretation] Format: boolean (see acl_group_enable). [internal type: boolean]  | 9 |
| The following attributes apply only to execution queues:   | 9 |
| checkpoint_min §   | 9 |
| Specifies the minimum interval of cpu time, in minutes, which is allowed between checkpoints of a job. If a user specifies a time less than this value, this value is used instead. Format: integer; default value: no minimum. [internal type: integer]   | 9 |
| resources_available  | 9 |
| The list of resource and amounts available to jobs running in this queue. The sum of the resource of each type used by all jobs running from this queue cannot exceed the total amount listed here. Advisory to the Scheduler, not enforced by the server. Format: "resources_available.resource_name=value", see qmgr(1B). [internal type: resource]  | 9 |
| kill_delay   | 9 |
| The amount of the time delay between the sending of SIGTERM and SIGKILL when a qdel command is issued against a running job. Format: integer seconds; default value: 2 seconds. [internal type: integer]   | 9 |
| max_user_run   | 9 |
| The maximum number of jobs owned by a single user that are allowed to be running from this queue at one time. This attribute is advisory to the Scheduler, it is not enforced by the server. Format: integer; default value: none. [internal type: integer]  | 9 |
| max_group_run  | 9 |
| The maximum number of jobs owned by any users in a single group that are allowed to be running from this queue at one time. This attribute is advisory to the Scheduler, it is not enforced by the server. Format: integer; default value: none. [internal type: integer]  | 9 |
| The following attributes apply only to routing queues:   | 9 |
| route_destinations   | 9 |
| The list of destinations to which jobs may be routed. Requires full manager privilege to set or alter. Format: comma separated strings of the form "queue_name[@server_host[:port]]". [default value: none, should be set to at least one valid destination] Example: "here,there@remote,test@remote:15501" [internal type: array of strings]  | 9 |
| alt_router   | 9 |
| If true, an site supplied, alternative job router function is used to determine the destination for routing jobs from this queue. Otherwise, the default, round-robin  | 9 |

|  |   |
|--|---|
| router is used. Requires full manager privilege to set or alter. Format: boolean (see acl_group_enable); default value: false. [internal type: boolean]  | 9 |
| route_held_jobs  | 9 |
| If true, jobs with a hold type set may be routed from this queue. If false, held jobs are not to be routed. Format: boolean (see acl_group_enable); default value: false. [internal type: boolean]   | 9 |
| route_waiting_jobs   | 9 |
| If true, jobs with a future execution_time attribute may be routed from this queue. If false, they are not to be routed. Format: boolean (see acl_group_enable); default value: false. [internal type: boolean]  | 9 |
| route_retry_time   | 9 |
| Time delay between route retries. Typically used when the network between servers is down. Format: integer seconds; default value: <b>PBS_NET_RETRY_TIME</b> (30 seconds). [internal type: integer]  | 9 |
| route_lifetime   | 9 |
| The maximum time a job is allowed to exist in a routing queue. If the job cannot be routed in this amount of time, the job is aborted. If unset or set to a value of zero (0), the lifetime is infinite. Format: integer seconds; default infinite. [internal type: integer]   | 9 |
| <br>   |   |
| 2.3.2. Queue Read-Only Attributes  | 9 |
| The following data items are read-only attributes of the queue. They are visible to but cannot be changed by clients.  | 9 |
| Items which apply to all types of queues are:  | 9 |
| total_jobs   | 9 |
| The number of jobs currently residing in the queue. [internal type: integer]   | 9 |
| state_count  | 9 |
| The total number of jobs currently residing in the queue in each state. [internal type: special, array of integers]  | 9 |
| These read-only attributes only apply to execution queues:   | 9 |
| resources_assigned   | 9 |
| The total amount of certain types of resources allocated to jobs running from this queue. [internal type: resource]  | 9 |
| hasnodes   | 9 |
| This attribute is set true if there are nodes associated with this queue. [internal type: boolean]   | 9 |
| <br>   |   |
| 2.4. <b>Batch Server Attributes</b>  | 9 |
| The following attributes apply to the server.  | 9 |
| <b>NAME</b>  | 9 |
| pbs_server_attributes - pbs server attributes  | 9 |
| <b>DESCRIPTION</b>   | 9 |
| Server attributes can be read by any client; privilege is not required. Most server attributes are alterable by a privileged client, run by a user with administrator or operator privilege. Certain attributes require the user to have full administrator privilege. The following is a list of the server attributes. | 9 |

|  |   |
|--|---|
| acl_host_enable  | 9 |
| Attribute which when true directs the server to use the <i>acl_hosts</i> access control lists. Requires full manager privilege to set or alter. Format: boolean, "TRUE", "True", "true", "Y", "y", "1", "FALSE", "False", "false", "N", "n", "0"; default value: false = disabled.   | 9 |
| acl_hosts  | 9 |
| List of hosts which may request services from this server. This list contains the network name of the hosts. Local requests, i.e. from the server's host itself, are always accepted even if the host is not included in the list. See section 10.1, Authorization, in the PBS External Reference Specification. Requires full manager privilege to set or alter. Format: "[+ -]hostname.domain[...]"; default value: all hosts.             | 9 |
| acl_resv_host_enable   | 9 |
| Attribute which when true directs the server to use the <i>acl_resv_hosts</i> access control list. Requires full manager privilege to set or alter. Format: boolean, "TRUE", "True", "true", "Y", "y", "1", "FALSE", "False", "false", "N", "n", "0"; default value: false = disabled.   | 9 |
| acl_resv_hosts   | 9 |
| List of hosts which may request reservation services from this server. This list contains the network name of the hosts. Local requests, i.e. from the server's host itself, are always accepted even if the host is not included in the list. See section 10.1, Authorization, in the PBS External Reference Specification. Requires full manager privilege to set or alter. Format: "[+ -]hostname.domain[...]"; default value: all hosts. | 9 |
| acl_resv_group_enable  | 9 |
| Attribute which when true directs the server to use the reservation group access control list <i>acl_resv_groups</i> . Requires full manager privilege to set or alter. Format: boolean, "TRUE", "True", "true", "Y", "y", "1", "FALSE", "False", "false", "N", "n", "0"; default value: false = disabled.   | 9 |
| acl_resv_groups  | 9 |
| List which allows or denies accepting reservations owned by members of the listed groups. The groups in the list are groups on the server host, not submitting hosts. See section 10.1, Authorization, in the PBS External Reference Specification. Format: "[+ -]group_name[...]"; default value: all groups allowed.   | 9 |
| acl_user_enable  | 9 |
| Attribute which when true directs the server to use the server level <i>acl_users</i> access list. Requires full manager privilege to set or alter. Format: boolean (see <i>acl_group_enable</i> ); default value: disabled.   | 9 |
| acl_users  | 9 |
| List of users allowed or denied the ability to make any requests of this server. See section 10.1, Authorization, in the PBS External Reference Specification. Requires full manager privilege to set or alter. Format: "[+ -]user[@host][...]"; default value: all users allowed.   | 9 |
| acl_resv_user_enable   | 9 |
| Attribute which when true directs the server to use the server level <i>acl_resv_users</i> access list. Requires full manager privilege to set or alter. Format: boolean (see <i>acl_group_enable</i> ); default value: disabled.  | 9 |
| acl_resv_users   | 9 |
| List of users allowed or denied the ability to make any reservation requests of this server. See section 10.1, Authorization, in the PBS External Reference Specification. Requires full manager privilege to set or alter. Format: "[+ -]user[@host][...]"; default value: all users allowed.   | 9 |

|   |   |
|---|---|
| acl_roots   | 9 |
| List of super users who may submit to and execute jobs at this server. If the job execution id would be zero (0), then the job owner, root@host, must be listed in this access control list or the job is rejected. Requires full manager privilege to set or alter. Format: "[+ -]user[@host][,...]"; default value: no root jobs allowed.   | 9 |
| comment   | 9 |
| A text string which may be set by the scheduler or other privileged client to provide information to the batch system users. Format: any string; default value: none.   | 9 |
| default_node  | 9 |
| A node specification to use if there is no other supplied specification. This attribute is only used by servers where a <i>nodes</i> file exist in the <i>server_priv</i> directory providing a list of nodes to the server. If the <i>nodes</i> file does not exist, this attribute is not set by default and is ignored if set. The default value allows for jobs to share a single node. Requires full manager privilege to set or alter. Format: a node specification string; default value: 1#shared.  | 9 |
| default_queue   | 9 |
| The queue which is the target queue when a request does not specify a queue name. Format: a queue name; default value: none, must be set to an existing queue.  | 9 |
| flatuid   | 9 |
| If set true, this boolean indicates that the user execution ID, UID, space is flat, constant, across all systems from which a user may submit a job to this server. Therefore, if the job will execute under the UID of the job owner, the server will not need to authorize execution with that UID. If the job is to execute under a user name supplied in the job user list, see -u option, then authorization will take place. Requires full manager privilege to set or alter. Default value: unset – authorization is required for all UIDs | 9 |
| log_events  | 9 |
| A bit string which specifies the type of events which are logged, see the section on Event Logging in chapter 3 of the ERS. Format: integer; default value: 511, all events.  | 9 |
| mail_uid  | 9 |
| The uid from which server generated mail is sent to users. Requires full manager privilege to set or alter. Format: integer uid; default value: 0 for root.   | 9 |
| managers  | 9 |
| List of users granted batch administrator privileges. Format: <i>user@host.sub.domain[,user@host.sub.domain...]</i> . The host, sub-domain, or domain name may be “wild carded” by the use of an “*” character, see the description of user access control lists in chapter 10.1.1 of the ERS. Requires full manager privilege to set or alter. Default value: root on the local host.  | 9 |
| max_running   | 9 |
| The maximum number of jobs allowed to be selected for execution at any given time. Advisory to the Scheduler, not enforced by the server. Format: integer.  | 9 |
| max_user_run  | 9 |
| The maximum number of jobs owned by a single user that are allowed to be running from this queue at one time. This attribute is advisory to the Scheduler, it is not enforced by the server. Format: integer; default value: none.  | 9 |
| max_group_run   | 9 |
| The maximum number of jobs owned by any users in a single group that are allowed to be running from this queue at one time. This attribute is advisory to the   | 9 |

|                                  |   |   |
|----------------------------------|---|---|
|                                  | Scheduler, it is not enforced by the server. Format: integer; default value: none.  | 9   |
| <code>node_pack</code>           | Controls how multiple processor nodes are allocated to jobs. If this attribute is set to true, jobs will be assigned to the multiple processor nodes with the fewest free processors. This packs jobs into the fewest possible nodes leaving multiple processor nodes free for jobs which need many processors on a node. If set to false, jobs will be scattered across nodes reducing conflicts over memory between jobs. If unset, the jobs are packed on nodes in the order that the nodes are declared to the server (in the nodes file). Requires full manager privilege to set or alter. Default value: unset – assigned to nodes as nodes in order that were declared.  | 9<br>9<br>9<br>9<br>9<br>9<br>9<br>9      |
| <code>operators</code>           | List of users granted batch operator privileges. Format of the list is identical with managers above. Requires full manager privilege to set or alter. Default value: root on the local host.   | 9<br>9<br>9                               |
| <code>query_other_jobs</code>    | The setting of this attribute controls if general users, other than the job owner, are allowed to query the status of or select the job. Format: boolean (see <code>acl_host_enable</code> ); Requires full manager privilege to set or alter. default value: false - users may not query or select jobs owned by other users.  | 9<br>9<br>9<br>9                          |
| <code>resources_available</code> | The list of resource and amounts available to jobs run by this server. The sum of the resource of each type used by all jobs running by this server cannot exceed the total amount listed here. Advisory to the Scheduler, not enforced by the server. Format: "resources_available.resource_name=value[,...]".   | 9<br>9<br>9<br>9                          |
| <code>resources_cost</code>      | The cost factors of various types of resources. These values are used in determining the order of releasing members of synchronous job sets, see the section on "Synchronize Job Starts." For the most part, these value are purely arbitrary and have meaning only in the relative values between systems. The "cost" of the resources requested by a job is the sum of the products of the various <code>resources_costs</code> and the amount of each resource requested by the job. It is not necessary to assign a cost for each possible resource, only those which the site wishes to be considered in synchronous job scheduling. Requires full manager privilege to set or alter. Format: "resources_cost.resource_name=value[,...]"; default value: none, cost of resource is not computed. | 9<br>9<br>9<br>9<br>9<br>9<br>9<br>9<br>9 |
| <code>resources_default</code>   | The list of default resource values that are set as limits for a job executing on this server when the job does not specify a limit, and there is no queue default. Format: "resources_default.resource_name=value[,...]"; default value: no limit.   | 9<br>9<br>9                               |
| <code>resources_max</code>       | The maximum amount of each resource which can be requested by a single job executing on this server if there is not a <code>resources_max</code> valued defined for the queue in which the job resides. Format: "resources_max.resource_name=value[,...]"; default value: infinite usage.   | 9<br>9<br>9<br>9                          |
| <code>scheduler_iteration</code> | The time, in seconds, between iterations of attempts by the batch server to schedule jobs. On each iteration, the server examines the available resources and runnable jobs to see if a job can be initiated. This examination also occurs whenever a running batch job terminates or a new job is placed in the queued state in an execution queue. Format: integer seconds; default value: 10 minutes, set by <b>PBS_SCHEDULE_CYCLE</b> in <code>server_limits.h</code> .   | 9<br>9<br>9<br>9<br>9<br>9                |

|   |   |
|---|---|
| scheduling  | 9 |
| Controls if the server will request job scheduling by the PBS job scheduler. If true, the scheduler will be called as required; if false, the scheduler will not be called and no job will be placed into execution unless the server is directed to do so by an operator or administrator. Setting or resetting this attribute to true results in an immediate call to the scheduler. For more information, see the section <b>Scheduler – Server Interaction</b> in the PBS Administrator Guide. Format: boolean (see <code>acl_host_enable</code> ); default value: value of <code>-a</code> option when server is invoked, if <code>-a</code> is not specified, the value is recovered from the prior server run. If it has never been set, the value is "false". | 9 |
| system_cost   | 9 |
| An arbitrary value factored into the resource cost of any job managed by this server for the purpose of selecting which member of synchronous set is released first, see <code>resources_cost</code> and section 3.2.2, "Synchronize Job Starts." Requires full manager privilege to set or alter. [default value: none, cost of resource is not computed]  | 9 |
| <b>Read Only Server Attributes</b>  | 9 |
| The following attributes are read-only, they are maintained by the server and cannot be changed by a client.  | 9 |
| resources_assigned  | 9 |
| The total amount of certain types of resources allocated to running jobs.   | 9 |
| server_name   | 9 |
| The name of the server which is the same as the host name. If the server is listening to a non-standard port, the port number is appended, with a colon, to the host name. For example: <i>host.domain:9999</i> .   | 9 |
| server_state  | 9 |
| The current state of the server:  | 9 |
| Active The server is running and will invoke the job scheduler as required to schedule jobs for execution.  | 9 |
| Idle The server is running but will not invoke the job scheduler.   | 9 |
| Scheduling  | 9 |
| The server is running and there is an outstanding request to the job scheduler.   | 9 |
| Terminating   | 9 |
| The server is terminating. No additional jobs will be scheduled.  | 9 |
| Terminating, Delayed  | 9 |
| The server is terminating in delayed mode. The server will not run any new jobs and will shutdown when the last currently executing job completes.  | 9 |
| state_count   | 9 |
| The total number of jobs managed by the server currently in each state.   | 9 |
| total_jobs  | 9 |
| The total number of jobs currently managed by the server.   | 9 |
| PBS_version   | 9 |
| The release version number of the server.   | 9 |
| <b>SEE ALSO</b>   | 9 |
| the PBS ERS, <code>qmgr(1B)</code>  | 9 |

|   |   |
|---|---|
| 2.4.1. Server Sub-Objects   | 9 |
| The following classes of objects are used by the server.  | 9 |
| queueServer manages a list of queue objects   | 9 |
| job Server manages lists of job objects.  | 9 |
| pbs_node  | 9 |
| Server manages lists of pbs_node objects.   | 9 |
| resc_resv   | 9 |
| Server manages lists of resc_resv (reservation) objects.  | 9 |
| <br>  |   |
| 2.5. Node Attributes  | 9 |
| The following attributes apply to pbs_node objects.   | 9 |
| <br>  |   |
| <b>NAME</b>   | 9 |
| pbs_node_attributes - pbs node attributes   | 9 |
| <br>  |   |
| <b>DESCRIPTION</b>  | 9 |
| <br>  |   |
| 2.5.1. Node Public Attributes   | 9 |
| Node public attributes are alterable on request by a client. The client must be acting for a user with administrator (manager) or operator privilege. Certain attributes require the user to have full administrator privilege before they can be modified.   | 9 |
| ntype This attribute defines the type of the node. Currently there are three types of nodes supported:  | 9 |
| time-shared   | 9 |
| Time-shared nodes are assumed to run multiple jobs and the placement of the jobs among the time-shared nodes is by the site policy defined in the Job Scheduler.  | 9 |
| cluster   | 9 |
| Cluster nodes are typically used for running parallel jobs. A cluster node is allocated to one or more jobs by the Server. Allocation is at the CPU level. A cluster node with one CPU declared can only be allocated to one job. A cluster node with N CPUs declared can be allocated one CPU to N jobs or N CPUs to one job are any combination there of. | 9 |
| globusA special pbs_mom is running to hand off jobs to the Globus distributed system.   | 9 |
| Requires full manager privilege to set or alter. Format: string, "time-shared", "cluster", "globus"; default value: cluster [internal type: short]  | 9 |
| state Shows or sets the state of the node. Certain state values, marked with an * in the following list, may be set by the manager or operator, the other states are set internally.  | 9 |
| free *Node is up and capable of accepting additional job(s).  | 9 |
| offline *   | 9 |
| Node has been marked by operator or administrator as unusable.  | 9 |
| downNode is not responding to queries from the Server.  | 9 |
| job-busy  | 9 |
| All CPUs on the node are allocated to jobs.   | 9 |



|  |   |
|--|---|
| job-sharing  | 9 |
| All CPUs on the node are allocated to jobs, at least one job is willing to share its allocated CPUs with other jobs that are also willing to share.  | 9 |
| job-exclusive  | 9 |
| The entire node has been exclusively allocated to one job at the job's request.  | 9 |
| busy   | 9 |
| The node is reporting a load average greater than the configured high water value.   | 9 |
| state-unknown  | 9 |
| The Server has never been able to contact the node. Either pbs_mom is not running on the node, the node hardware is down, or there is a network problem.   | 9 |
| Format: string, one of the above states. [internal type: short]  | 9 |
| properties   | 9 |
| Properties are arbitrary strings assigned to the node by the administrator. Users may request nodes by specifying one or more properties. The notion of properties may be used to group nodes by certain characteristics. Requires full manager privilege to set or alter. Format: comma separated list of alphanumeric (first character must be alphabetic) strings; default value: none. [internal type: array of strings] | 9 |
| max_running  | 9 |
| The maximum number of jobs allowed to be run on this node at any given time. This attribute is advisory to the Scheduler, it is not enforced by the server. Format: integer. [internal type: integer]  | 9 |
| max_user_run   | 9 |
| The maximum number of jobs owned by a single user that are allowed to be run on this node at one time. This attribute is advisory to the Scheduler, it is not enforced by the server. Format: integer; default value: none. [internal type: integer]   | 9 |
| max_group_run  | 9 |
| The maximum number of jobs owned by any users in a single group that are allowed to be run on this node at one time. This attribute is advisory to the Scheduler, it is not enforced by the server. Format: integer; default value: none. [internal type: integer]   | 9 |
| resources_available  | 9 |
| The list of resource and amounts available on this node. If not explicitly set, the amount shown is that reported by the pbs_mom running on the node. Currently, only the ncpus number will be retained across Server restarts. Format: "resources_available.resource_name=value", see qmgr(1B). [internal type: resource]   | 9 |
| queue  | 9 |
| The queue with which this node is associated. If set, only jobs in that queue may run on this node. If not set, any job in a queue without associated nodes may run on this node. Requires full manager privilege to set or alter. Format: "queue name"; default value: none. [internal type: string]  | 9 |
| <br>   |   |
| 2.5.2. Node Read-Only Attributes   | 9 |
| The following attributes are read only:  | 9 |
| jobs   | 9 |
| List of jobs running on the node. This only applies to cluster nodes. This attribute is read-only. Format: "#/jobid,...", where # represents the number of the processor. [internal type: string]  | 9 |

|  |   |
|--|---|
| reservations   | 9 |
| List of advance reservations pending on the node. This attribute is read-only. Format: "#/reservation id,..." [internal type: list of strings]   | 9 |
| resources_assigned   | 9 |
| The total amount of certain types of resources allocated to jobs running on this node. This attribute is read-only. [internal type: resource]  | 9 |
| <br>   |   |
| <b>2.6. Reservation Attributes</b>   | 9 |
| The following attributes apply to resc_resv (resources reservation) objects.   | 9 |
| <br>   |   |
| <b>NAME</b>  | 9 |
| pbs_reservation_attributes - pbs advance resources reservation attributes  | 9 |
| <br>   |   |
| <b>DESCRIPTION</b>   | 9 |
| Listed below are the "public" attributes available on an advance resources reservation entity. Attributes which are not directly set through user action (usually via parameters that are part of the resources reservation request) may still get assigned a default value by the server to which the resources reservation is submitted, while other ones of the object's attributes are not used and so receive no value. | 9 |
| <br>   |   |
| Account_Name   | 9 |
| Used for accounting on some hosts.   | 9 |
| Format: string.  | 9 |
| Default value: none.   | 9 |
| <br>   |   |
| group_list   | 9 |
| A list of <i>group_name@hosts</i> . It is used to determine what group name to assign to a resources reservation located on a given server host. When a resources reservation request is submitted to a particular server host, the server on that host selects a group name from this list according to an ordered set of rules:  | 9 |
| 1. Select from the list that group name for which the associated host name matches the name of the submitting host.  | 9 |
| 2. Select from the list that group name which has no associated host name, i.e. the "wild card name."  | 9 |
| 3. Use the login group for the name in the euser attribute. The name in the euser attribute is itself determined by the server from an ordered set of rules.   | 9 |
| Format: "group_name[@host][,group_name[@host]...]".  | 9 |
| <br>   |   |
| Reserve_Name   | 9 |
| The name assigned to the resources reservation during submission via the <b>pbs_rsub</b> or the <b>qsub</b> (in the case of a reservation job) command.  | 9 |
| Format: string up to 15 characters, where the first character is alphabetic.   | 9 |
| Default value: none.   | 9 |
| <br>   |   |
| Mail_Points  | 9 |
| Identifies which of the reservation state transition events cause mail notification to be sent by the server to various receiving parties.   | 9 |
| Format: string consisting of one or more letters (abce) with 'c' meaning notify when scheduler confirms, 'b' notify when reservation period begins, 'e' notify when reservation period ends, 'a' notify when reservation is terminated.  | 9 |
| Default value: "ac", send mail on confirmation and on abort/delete events only.  | 9 |
| <br>   |   |
| Mail_Users   | 9 |
| The set of users to whom mail may be sent when the reservation experiences various changes in state.   | 9 |

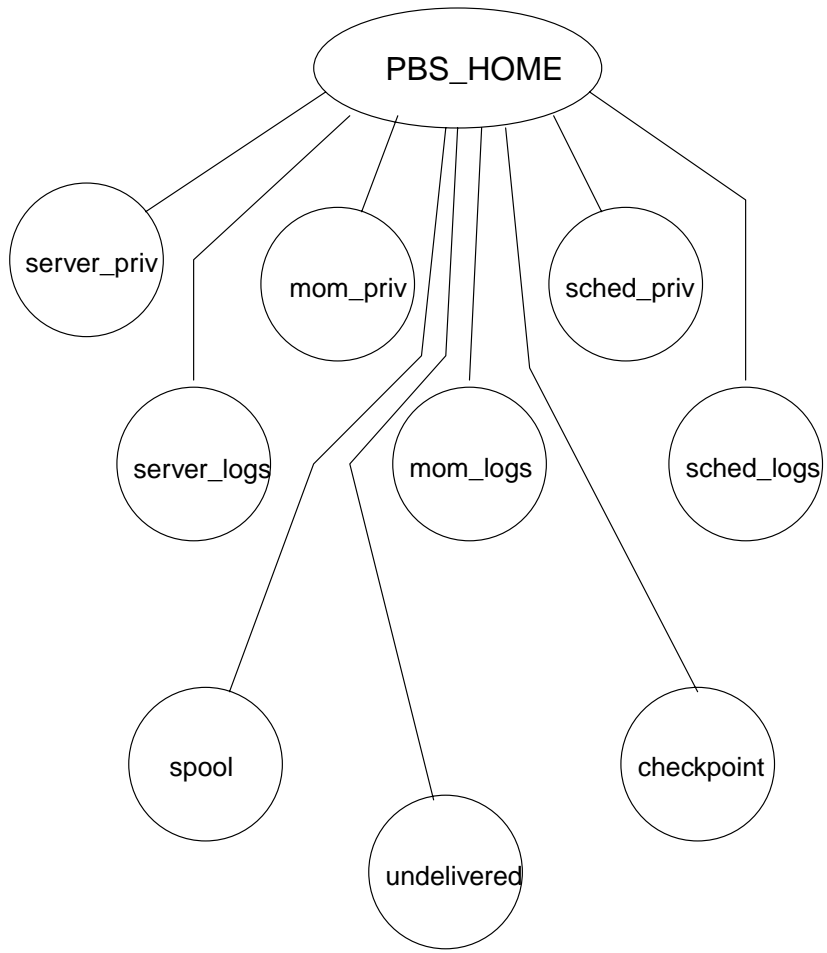
|   |   |
|---|---|
| Format: string of the form, user@host[,user@host].... .   | 9 |
| Default value: reservation owner only.  | 9 |
| <b>Priority</b>   | 9 |
| The reservation scheduling priority assigned by the user. This attribute is yet to be implemented.  | 9 |
| Format: string of the form, [+ -]nnnnn .  | 9 |
| Default value: none.  | 9 |
| <b>Resource_List</b>  | 9 |
| The list of resources requested by the reservation. This list is a set of <i>name=value</i> strings. The actual meaning of <i>name</i> and <i>value</i> is server dependent. The values specified for the various requested resources serve as limits on the aggregate of those jobs that are, at any moment, running under the reservation. Default value: none  | 9 |
| <b>User_List</b>  | 9 |
| The list of <i>user@hosts</i> which determines what user name gets assigned to the reservation by the server to which it is sent. [internal type: array of strings] The server will use the following ordered set of rules to decide a user name to associate with the reservation:   | 9 |
| 1. Select that user name from the list for which the associated host name matches the name of the server host.  | 9 |
| 2. Select from the list that user name which has no associated host name, the "wild card name."   | 9 |
| 3. Use the Reserve_Owner name as the user name.   | 9 |
| Default value: Reserve_Owner name.  | 9 |
| <b>Variable_List</b>  | 9 |
| A list of PBS environment variables passed with the <i>SubmitResv</i> batch request.  | 9 |
| Format: string of the form, name=value[,name=value]... .  | 9 |
| <br><b>Privileged Resource Reservation Attributes</b>   | 9 |
| The following attributes require system, manager, or operator privilege to set. Their visibility to a client can depends on privilege as noted.   | 9 |
| <br><b>2.6.1. Resources Reservation Read-Only Attributes</b>  | 9 |
| These next set of attributes are read-only. They are established by the server, are visible to the client, but can not be set by a client. Certain ones are only visible to "privileged" clients (those clients being run by the batch administrator).  | 9 |
| <b>reserve_type</b>   | 9 |
| Indicates the type of resources reservation. A value of (2) means this is a general reservation, (3) means this is a reservation job.   | 9 |
| <b>reserve_ID</b>   | 9 |
| The resources reservation's identifier. This is a string in one of two formats depending on the type of resources reservation. For a reservation job the the format of this id string is the same as that for a job, namely, [sequence_number.server_name@server]. In the case of a general reservation of resources the format is [Rsequence_number.server_name@server]. The pieces of this identifier are as follows: sequence_number is a positive sequential integer assigned by the server to which first created the reservation, server_name is the name of the host where the reservation was created, @server is the name of the server host on which the reservation currently resides. | 9 |

|  |   |
|--|---|
| reserve_start  | 9 |
| The time at which resources requested by the reservation can start to be used.   | 9 |
| reserve_end  | 9 |
| The time at which resources requested by the reservation can no longer be consumed. Also, the time when the resources reservation (reservation job) will begin to be automatically removed from the system.  | 9 |
| reserve_duration   | 9 |
| A contiguous amount of time. During this period of time resources requested in the reservation (reservation job) are to be available to satisfy jobs submitted to run in that reservation.   | 9 |
| resv_nodes   | 9 |
| The value of this attribute is the collection of nodes that will be selected by the PBS system to solve the nodes specification that is part of the resources reservation (reservation job) request.   | 9 |
| Format: a contiguous '+'-separated string of specific nodes specifications.  | 9 |
| Authorized_Users   | 9 |
| The list of those users that are given (denied) the right to submit jobs to the queue instantiated to service the confirmed resources reservation request. This list is used as the user_acl (access control list) for the queue. The list is currently not changeable by the owner of the reservation.  | 9 |
| Format: [+ -]user[hostname.domain],..., [+ -]... where, '-' means "deny" and   | 9 |
| Default value: only the reservation owner is allowed to submit to the queue.   | 9 |
| Authorized_Groups  | 9 |
| List of groups which allows (denies) the right of users belonging to those groups to enqueue jobs in the queue associated with the reservation. This list is used as the group_acl (access control list) for the instantiated queue.   | 9 |
| Format: [+ -]group_name,..., [+ -]group_name] (see Authorized_Users).  | 9 |
| Default value: owner's login group.  | 9 |
| ctime  | 9 |
| The time that the resources reservation object was created.  | 9 |
| egroup   | 9 |
| Attribute egroup is set to a name that gets determined by the server to which the resources reservation is sent. Only the batch administrator has availability to this attribute (see group_list.)   | 9 |
| euser  | 9 |
| Attribute euser is set to a user name that gets determined by the server to which the resources reservation is sent. Only the batch administrator has availability to this attribute.  | 9 |
| hashname   | 9 |
| The name used as a basename for the resources reservation file for this reservation. This file is part of the server's database and used in recovery of the resources reservation whenever the PBS server is coming back up. This attribute is available only to the batch administrator.  | 9 |
| Interactive  | 9 |
| This attribute is set to a non-zero value if the client that submitted the resources reservation request is willing to wait for the scheduler to confirm the reservation. The value of this attribute is the number of seconds that the client is willing to wait for confirmation. A positive number is interpreted as being the number of seconds that the client is willing to wait to get back confirmation of the request, after which confirmation of the reservation will be done by querying via <b>pbs_rstat</b> . If the value of this attribute is negative, it is interpreted to mean that the client is willing to wait that many seconds (positive) for confirmation by the scheduler and, if confirmation has not occurred in that period of time the submitted resources | 9 |

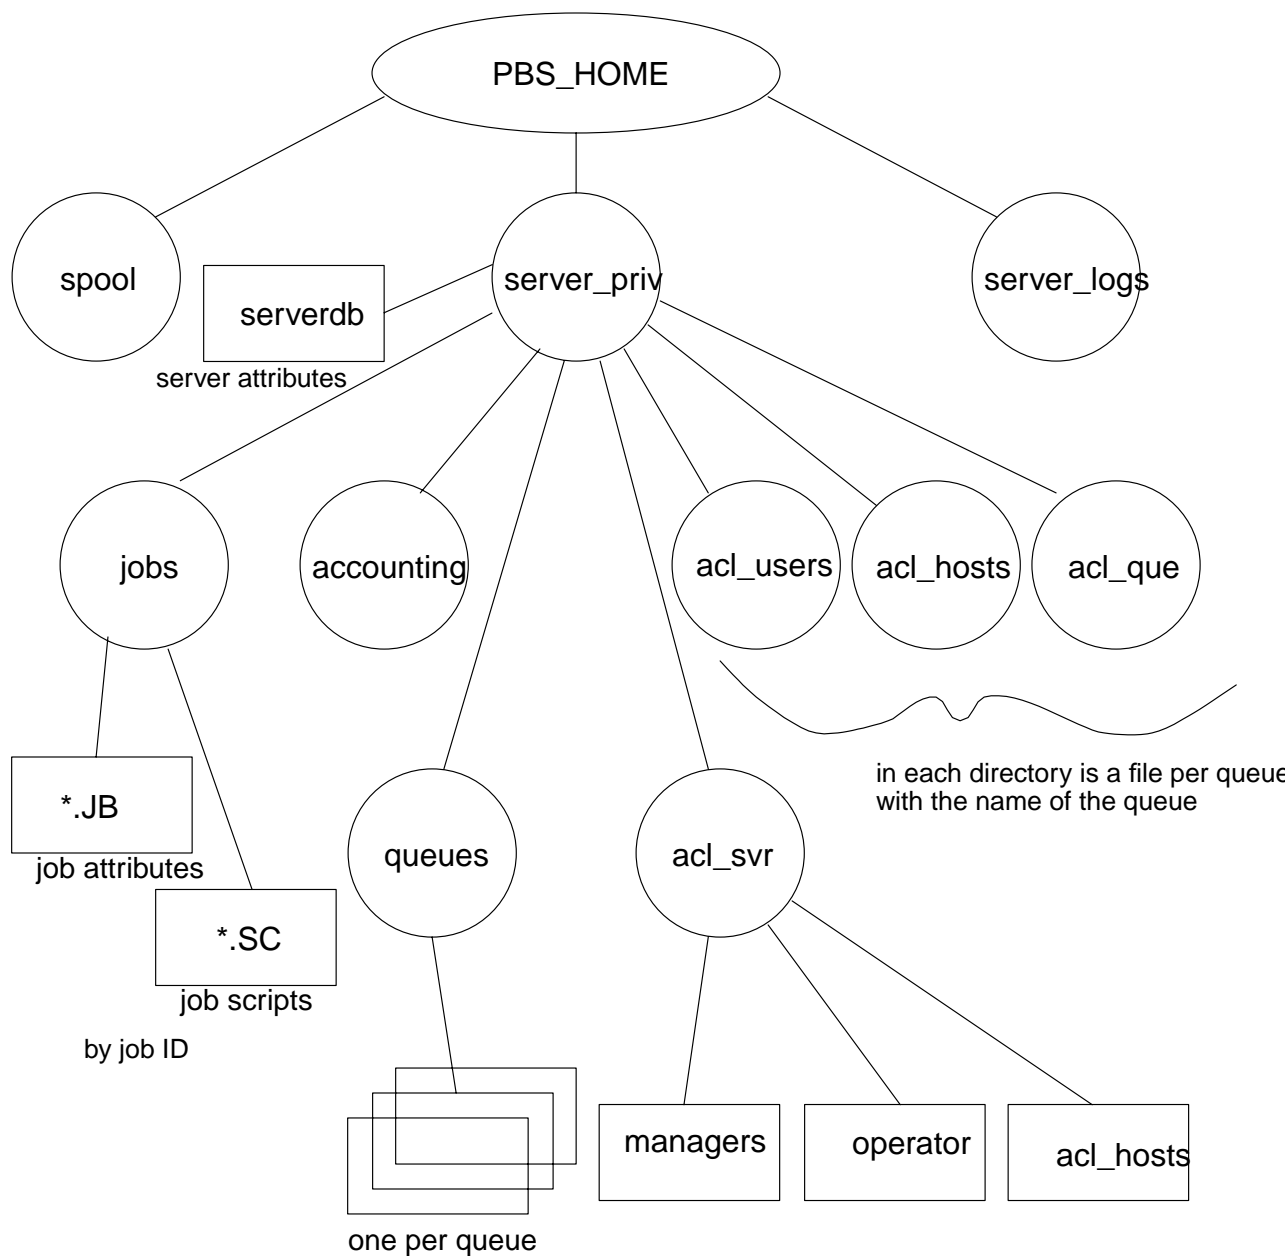
|   |   |
|---|---|
| reservation is to be automatically deleted from the system.   | 9 |
| <b>Reserve_Owner</b>  | 9 |
| The login name on the submitting host of the user who submitted the resources reservation request.  | 9 |
| <b>reserve_state</b>  | 9 |
| The state of the resources reservation. The various states are:   | 9 |
| <b>RESV_UNCONFIRMED</b>   | 9 |
| resources reservation (reservation job) request is awaiting scheduler confirmation.   | 9 |
| <b>RESV_CONFIRMED</b>   | 9 |
| scheduler has confirmed the resources reservation (reservation job) request. For a general resources reservation, a reservation queue has gotten instantiated, jobs and reservation jobs can now be enqueued into this queue and later, during the reservation period, be selected by the scheduler to run. | 9 |
| <b>RESV_WAIT</b>  | 9 |
| confirmed reservation job is waiting arrival of the reservation window.   | 9 |
| <b>RESV_TIME_TO_RUN</b>   | 9 |
| scheduler signaled by server upon the arrival of the reservation period   | 9 |
| <b>RESV_RUNNING</b>   | 9 |
| resources reservation period has started  | 9 |
| <b>RESV_FINISHED</b>  | 9 |
| resources reservation period is now finished  | 9 |
| <b>RESV_BEING_DELETED</b>   | 9 |
| process of deleting the resources reservation (reservation job) commencing.   | 9 |
| <b>RESV_DELETED</b>   | 9 |
| resources reservation is now being deleted, any jobs (reservation jobs) that belonged to the reservation are already deleted.   | 9 |
| <b>RESV_DELETING_JOBS</b>   | 9 |
| jobs belonging to the resources reservation are being deleted prior to deleting the resources reservation itself.   | 9 |
| <b>Queue</b>  | 9 |
| The name of the instantiated execution (a.k.a "reservation") queue. Jobs that are to use resources belonging to this reservation need to be submitted to this queue.  | 9 |
| <b>substate</b>   | 9 |
| A numerical indicator of the substate of the resources reservation. The substate is used internally by the PBS server in managing the reservation. The attribute is visible to privileged clients, such as the scheduler.   | 9 |
| Format: interger. The values of substate are defined in the header file reservation.h   | 9 |
| <b>SEE ALSO</b>   | 9 |
| PBS ERS, qsub(1B), qalter(1B)   | 9 |

## 2.7. PBS Files 9

The PBS subsystem maintains its files under the directory {PBS\_DIR} which is usually set to /usr/spool/PBS, see figure 2-1 for the layout directories used by PBS. There is a subdirectory, {PBS\_DIR} /server\_priv for private files accessible only by the server, see figure 2-2 used by the Server daemon for its private files. The Job Executor, MOM, has a similar subdirectory mom\_priv. 9



**Figure 2-1: PBS Directory Structure**



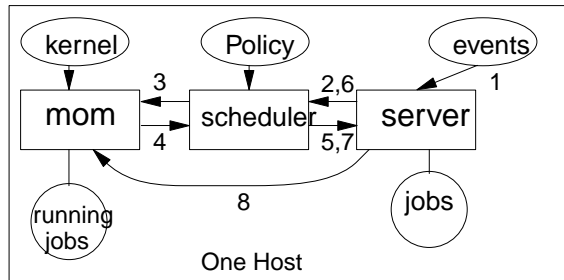
**Figure 2-3: PBS Server Files**

**2.8. Job Selection / Scheduling** 9

Job selection or job scheduling is the process of picking which eligible job is placed into execution or initiated. To be eligible, the job must reside in an execution queue and be in the **queued** state. 9

To maximize flexibility in implementing site policy, **PBS** provides a separate program as the selection process. This process operates on the principle of evaluating the eligible jobs according to a program written either in a yacc/lex based procedural language or in Tcl. This process, the **PBS Scheduler**, `pbs_sched`, communicates with the PBS Server via a socket based IPC, using the standard PBS API. This provides the capabili- 9

ty of having the Scheduler and the Server reside on different hosts. 9  
 The Scheduler communicates with another process called the **Machine Oriented** 9  
**Miniserver**, pbs\_mom, in its role as a resource manager to obtain information about 9  
 the loading of the host system. Loading information may contain details of memory us- 9  
 age, cpu load average, and more. This information, taken as input to the scheduling 9  
 program can be used to control job scheduling. The relationship between PBS, the 9  
 Scheduler, and the Machine Oriented Miniserver can be seen in figures 2-4 and 2-5. 9

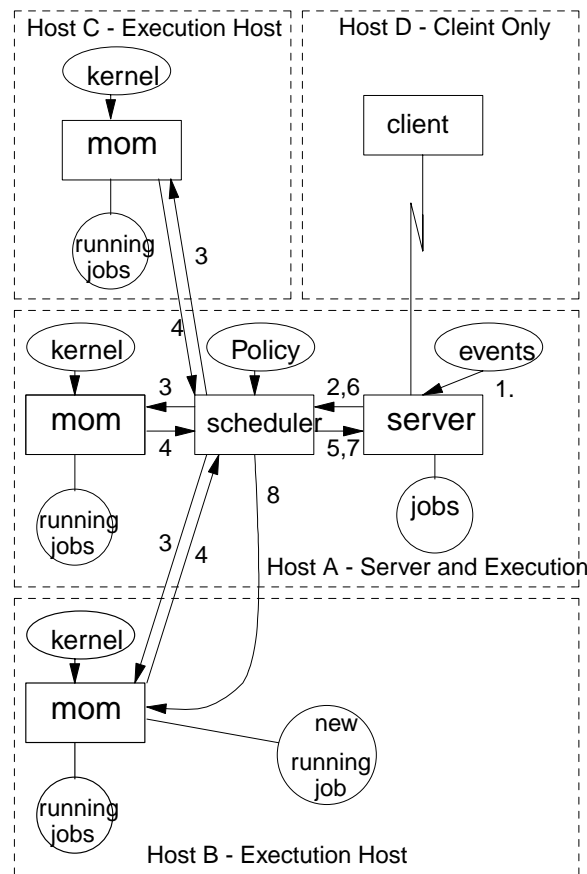


1. Event tells Server to initiate a scheduling cycle.
2. Server sends scheduling command to Scheduler.
3. Scheduler requests resource info from MOM.
4. MOM returns requested info.
5. Scheduler requests job info from server.
6. Server sends job status info to scheduler.
- Scheduler makes policy decision to run job.
7. Scheduler sends run request to server.
8. Server sends job to MOM to run.

**Figure 2-4: Batch Scheduling on a Single Host**

The concept of the Scheduler and Machine Oriented Miniserver can be extended to in- 9  
 clude multiple hosts. A single Scheduler would provide scheduling for one or more PBS 9  
 servers. The Scheduler would talk with a Machine Oriented Miniserver on each host 9  
 for which it had scheduling responsibility. 9





**Figure 2-5: Batch Scheduling on Multiple Hosts**

The procedural or Tcl script which drives the Scheduler is a plain text file. The administrator may update the script in the file or direct the scheduler to use a different file. 9  
 The script may be written to consider time and queue information, so different procedures could apply at different times and to different queues. The script may also state when the conditions are such that a job should be started, deleted, or suspended. 9

**2.9. General Identifiers** 9

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 sub-sections and to the general rules described in the next paragraph. 9

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. 9

**2.9.1. Account String** 9

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 9

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. 9  
9  
9

2.9.2. Attribute Name 9

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. 9  
9  
9  
9

2.9.3. Destination Identifiers 9

A destination identifier is a string used to specify a particular destination. The identifier may be specified in one of three forms: 9  
9

```
queue@server_name
queue
@server_name
```

where: 9

`queue` 9  
is an ASCII character string of up to 15 characters. Valid characters are alphanumeric, 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. 9  
9

`server_name` 9  
is a string identifying a server; see `server_name`, section 2.7.9. If `server_name` is omitted, the default server is assumed. 9  
9

2.9.4. Default Server 9

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 by (b) the server name in the file specified by the `$(PBS_DEFAULT_FILE)` build parameter in the local.mk file. 9  
9  
9

2.9.5. Host Name 9

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. 9  
9  
9  
9  
9  
9  
9  
9

2.9.6. Job Identifiers 9

When the term *job identifier* is used, the identifier is specified as: 9

```
sequence_number[.server_name][@server]
```

The `sequence_number` is the number supplied by the server when the job was submitted. 9

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. 9

`@server` specifies the current location of the job. See the definition of default server in section 2.7.4 and the section 5.1.2, entitled "Directing Requests to Correct Server." 9

When the term *fully qualified job identifier* is used, the identifier is specified as: 9

```
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. 9

### 2.9.7. Job Name 9

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. 9

### 2.9.8. Resource Name 9

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 in section 3.4.3, “Types of Resources”. 9

### 2.9.9. Server Name 9

*Server Name* is an ASCII character string of the form: 9

```
basic_server_name[:port]
```

The string identifies a batch server. Basic server names are identical to host names (see section 2.7.5). 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. 9

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. 9

### 2.9.10. User Name 9

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. 9

|  |   |
|--|---|
| <b>3. Batch Server Functions</b>   | 9 |
| A batch server provides services in one of two ways:   | 9 |
| - The server provides a service at the request of a client.  | 9 |
| - The server provides a <i>deferred service</i> 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.   | 9 |
| <br>   |   |
| <b>3.1. Client Service Requests</b>  | 9 |
| 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.   | 9 |
| The server is required to respond to all requests it receives. Those requests that cannot be successfully completed, are <i>rejected</i> . The reason for the rejection is returned in the reply to the client.  | 9 |
| The following subsections describe the services provided by a batch server in response to a request from a client. The requests are grouped in the following sub-sections by the type of object affected by the request: server, queue, job, or resource.  | 9 |
| <br>   |   |
| <b>3.1.1. Server Management</b>  | 9 |
| 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.  | 9 |
| <br>   |   |
| <b>3.1.1.1. Manage Request</b>   | 9 |
| The <i>Manage</i> request supports the <b>qmgr(8)</b> 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.   | 9 |
| <br>   |   |
| <b>3.1.1.2. Server Status Request</b>  | 9 |
| The status of the server may be requested with a <i>Server Status</i> request.   | 9 |
| The batch server will reject the request if any of the following conditions are true:  | 9 |
| - The user of the client is not authorized to query the status of the server.  | 9 |
| If the request is accepted, the server will return a <i>Server Status Reply</i> . See the <b>qstat</b> command and the Data Exchange Format description for details of which server attributes are returned to the client.   | 9 |
| <br>   |   |
| <b>3.1.1.3. Start Up</b>   | 9 |
| 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.   | 9 |
| The server is started by a <b>pbs_server</b> 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 <b>running</b> state when the server previously shut down is dictated by the start up mode, see the description of the pbs_server(8) command. | 9 |

|   |   |
|---|---|
| 3.1.1.4. Shut Down  | 9 |
| 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 <i>Server Shutdown</i> request.  | 9 |
| 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 <b>qterm</b> command.   | 9 |
| 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. | 9 |
| If the server receives either a <b>SIGTERM</b> or a <b>SIGSHUTDN</b> signal, the server will act as if it had received a shut down immediate request.   | 9 |
| <br>  |   |
| 3.1.2. Queue Management   | 9 |
| 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.   | 9 |
| <br>  |   |
| 3.1.2.1. Queue Status Request   | 9 |
| The status of a queue at the server may be requested with a <i>Queue Status</i> request.  | 9 |
| The batch server will reject the request if any of the following conditions are true:   | 9 |
| - The user of the client is not authorized to query the status of the designated queue.   | 9 |
| - The designated queue does not exist on the server.  | 9 |
| If the request does not specify a queue, status of all the queues at the server will be returned.   | 9 |
| When the request is accepted, the server will return a <i>Queue Status Reply</i> . See the <b>qs-tat</b> command and the Data Exchange Format description for details of which queue attributes are returned to the client.   | 9 |
| <br>  |   |
| 3.1.3. Job Management   | 9 |
| 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.   | 9 |
| <br>  |   |
| 3.1.3.1. Abort Request  | 9 |
| [The Abort Request has been aborted (deleted).]   | 9 |
| <br>  |   |
| 3.1.3.2. Commit Request   | 9 |
| The <i>Commit</i> 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.   | 9 |

3.1.3.3. Delete Job Request 9

A *Delete Job* request asks a server to remove a job from the queue in which it exists and not place it elsewhere. 9

The batch server will reject a Delete Job Request if any of the following conditions are true: 9

- The user of the client is not authorized to delete the designated job. 9
- The designated job is not owned by the server. 9
- The designated job is not in an eligible state. Eligible states are **queued**, **held**, **waiting**, **running**, and **transiting**. 9

If the job is in the **running** state, the server 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 server will send a **SIGKILL** signal to the job process group. 9

The job is then placed into the **exiting** state. 9

3.1.3.4. Hold Job Request 9

A client can request that one or more holds be applied to a job. 9

The batch server will reject a *Hold Job* request if any of the following conditions are true: 9

- The user of the client is not authorized to add any of the specified holds. 9
- The batch server does not manage the specified job. 9

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. 9

If the job is in the **queued** or **waiting** state, it is placed in the **held** state. 9

If the job is in **running** state, then the following additional actions are taken: If checkpoint/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 the *expedite modifier* to the hold type is specified, the sched\_hint job attribute will be set to one (1) to indicate to the scheduler that this job should be placed into execution as soon as possible. If the *expedite modifier* is not specified, the job will be rescheduled for execution as if it had been re-submitted. 9

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. 9

3.1.3.5. Queue Job Request 9

The *Queue Job* sub-request is part of the Queue Job complex request. This sub-request identifies the destination for the job and, in effect, asks the receiving server to fork off a child to process the remaining sub-requests. It also passes the public and private job attributes to the receiving server. 9

3.1.3.6. Job Credential Request 9

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 in section 10.2. 9

|   |   |
|---|---|
| 3.1.3.7. Job Script Request   | 9 |
| The <i>Job Script</i> 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. | 9 |
| 3.1.3.8. Locate Job Request   | 9 |
| 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 <i>Locate Job</i> request, it returns a <i>Locate Reply</i> .   | 9 |
| The request will be rejected if any of the following conditions are true:   | 9 |
| - The server does not own (manage) the job, and   | 9 |
| - The server did not create the job.  | 9 |
| - The server is not maintaining a record of the current location of the job.  | 9 |
| 3.1.3.9. Message Job Request  | 9 |
| 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.  | 9 |
| The batch server will reject a <i>Message Job</i> request if any of the following conditions are true:  | 9 |
| - The designated job is not in the running state.   | 9 |
| - The user of the client is not authorized to post a message to the designated job.   | 9 |
| - The designated job is not owned by the server.  | 9 |
| When the server accepts the <i>Message Job</i> request, it 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.   | 9 |
| 3.1.3.10. Modify Job Request  | 9 |
| A batch client makes a <i>Modify Job</i> 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:   | 9 |
| - The user of the client is not authorized to make the requested modification to the job.   | 9 |
| - The designated job is not owned by the server.  | 9 |
| - The requested modification is inconsistent with the state of the job. This is detailed later in this subsection.  | 9 |
| - A requested resource change would exceed the limits of the queue or server.   | 9 |
| - An unrecognized resource is requested for a job in an execution queue.  | 9 |
| 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.   | 9 |
| The following table indicates which attributes are alterable in each state.   | 9 |

| Attribute       | q | h | w | r | t | e |
|-----------------|---|---|---|---|---|---|
| Account_Name    | x | x | x |   |   |   |
| Checkpoint      | x | x | x | 1 |   |   |
| depend          | x | x | x |   |   |   |
| Error_Path      | x | x | x | 1 |   |   |
| Execution_Time  | x | x | x | 1 |   |   |
| group_list      | x | x | x |   |   |   |
| Hold_Types      | x | x | x | 1 |   |   |
| Job_Name        | x | x | x | x |   |   |
| Join_Path       | x | x | x |   |   |   |
| Keep_Files      | x | x | x |   |   |   |
| Mail_Points     | x | x | x | x |   |   |
| Mail_Users      | x | x | x | x |   |   |
| Output_Path     | x | x | x |   |   |   |
| Priority        | x | x | x |   |   |   |
| Rerunable       | x | x | x | x |   |   |
| Resource_List   | x | x | x | 2 |   |   |
| Shell_Path_List | x | x | x | 1 |   |   |
| stagein         | x | x | x |   |   |   |
| stageout        | x | x | x |   |   |   |
| User_List       | x | x | x |   |   |   |
| Variable_List   |   |   |   |   |   |   |

Notes:

1. May be altered with qalter, but changes will not take effect until job is re-queued. Use of qhold produces special processing.
2. Only certain resources limits may be changed. Those resource limits may be lowered by job owner. Increasing those limits requires special privilege.

3.1.3.11. Move Job Request 9

A client can request a server to move a job to a new destination. 9

The batch server will reject a *Move Job* Request if any of the following conditions are true: 9

- The user of the client is not authorized to remove the designated job from the queue in which the job resides. 9
- The user of the client is not authorized to submit a job to the new destination. 9
- The designated job is not owned by the server. 9
- The designated job is not in the **queued**, **held**, or **waiting** state. 9
- The new destination is disabled. 9
- The new destination is inaccessible. 9

When the server accepts a Move Job request, it will 9

- Queue the designated job at the new destination. 9
- Remove the job from the current queue. 9

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. 9

The server will insure that a job is neither lost nor duplicated. 9



### 3.1.3.12. Queue Job Request 9

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. 9

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. 9

The batch server will reject a Queue Job Request if any of the following conditions are true: 9

- The client is not authorized to create a job in the target queue. 9
- The target queue does not exist at the server. 9
- The target queue is not enabled. 9
- The target queue is an execution queue and a resource requirement of the job exceeds the limits set upon the queue. 9
- The target queue is an execution queue and an unrecognized resource is requested by the job. 9
- The target queue is an execution queue, the batch server does not support checkpoint, and the value of the Checkpoint attribute of the job is neither the single character "n" nor the single character "u". 9
- The job requires access to a user identifier that the client is not authorized to access. 9

When a job is placed in a execution queue, it is placed in the **queued** state unless one of the following conditions applies: 9

- The job has an execution\_time attribute that specifies a time in the future and the Hold\_Types attribute has value of {NONE}. Then the job is placed in the **waiting** state. 9
- The job has a Hold\_Types attribute with a value other than {NONE}. The job is placed in the **held** state. 9

When a job is placed in a routing queue, It is placed in the **queued** state unless one of the following conditions applies: 9

- The job has an execution\_time attribute that specifies a time in the future, the Hold\_Types attribute has value of {NONE}, and the route\_waiting\_jobs queue attribute is {FALSE}. Then the job is placed in the **waiting** state. 9
- The job has a Hold\_Types attribute with a value other than {NONE} and the route\_held\_jobs queue attribute is {FALSE}. Then the job is placed in the **held** state. 9

A batch server that accepts a Queue Job Request for a *new* job will 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. 9

A batch server that accepts a Queue Job Request for a *new* job will 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. 9

A batch server that accepts a Queue Job Request for a *new* job will 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. 9

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. 9

|   |   |
|---|---|
| 3.1.3.13. Release Job Request   | 9 |
| A client can request that one or more holds be removed from a job.  | 9 |
| A batch server rejects a <i>Release Job</i> request if any of the following conditions are true:  | 9 |
| - The user of the client is not authorized to add (remove) any of the specified holds.  | 9 |
| - The batch server does not manage the specified job.   | 9 |
| 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.   | 9 |
| If the job is in the <b>held</b> state and all holds have been removed, the job is placed in the <b>waiting</b> state if the Execution_Time attribute specifies a time in the future. Otherwise the job is placed in the <b>queued</b> state.   | 9 |
| <br>  |   |
| 3.1.3.14. Rerun Job Request   | 9 |
| 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. Unless the Hold_Types attribute is not {NONE}, the job is eligible to be re-scheduled for execution.  | 9 |
| The server will reject the <i>Rerun Job</i> request if any of the following conditions are true:  | 9 |
| - The user of the client is not authorized to rerun the designated job.   | 9 |
| - The Rerunable attribute of the job has the value {FALSE}.   | 9 |
| - The job is not in the running state.  | 9 |
| The server does not own the job.  | 9 |
| When the server accepts the Rerun Job request, it performs the following actions:   | 9 |
| - Send a <b>SIGKILL</b> signal to the session (process) group of the job.   | 9 |
| - 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 <b>held</b> state. If the execution_time attribute is a future time, the job will be placed in the <b>waiting</b> state. Otherwise, the job is placed in the <b>queued</b> state.  | 9 |
| <br>  |   |
| 3.1.3.15. Run Job   | 9 |
| The "Run Job" request directs the server to place the specified job into immediate execution. The request is issued by a <b>qrun</b> operator command and by the PBS Job Scheduler.   | 9 |
| <br>  |   |
| 3.1.3.16. Select Jobs Request   | 9 |
| 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 <i>Select Jobs</i> 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. | 9 |
| The server will reject the request if any of the following conditions are true:   | 9 |
| - The queue portion of a specified destination does not exist on the server.  | 9 |
| When the request is accepted, the server will return a <i>Select Reply</i> containing a list of zero or more jobs that met the selection criteria.  | 9 |
| <br>  |   |
| 3.1.3.17. Signal Job Request  | 9 |

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. 9

The batch server will reject a Signal Job Request if any of the following conditions are true: 9

- The user of the client is not authorized to signal the job. 9
- The job is not in the running state, except for the special signal “resume” when the job must be in the Suspended state. 9
- The server does not own the designated job. 9
- The requested signal is not supported by the host operating system. (The kill system call returns [EINVAL].) 9

When the server accepts a request to signal a job, it will send the signal requested by the client to the all processes in the job's session. 9

### 3.1.3.18. Status Job Request 9

The status of a job or set of jobs at a destination may be requested with a *Status Job* request. 9

The batch server will reject a Status Job Request if any of the following conditions are true: 9

- The user of the client is not authorized to query the status of the designated job. 9
- The designated job is not owned by the server. 9

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. 9

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. 9

## 3.2. Server to Server Requests 9

Server to Server requests are a special category of client requests. They are only issued to a server by another server. 9

### 3.2.1. Track Job Request 9

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. 9

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. 9

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. 9

### 3.2.2. Synchronize Job Starts 9

**PBS** provides for synchronizing the initiation of jobs across hosts. This is done to support distributing processing. 9

Author note: 9

There are several approaches that could be taken to solve this requirement, 9

none of them simple and straightforward. The best approach for synchronization of jobs would be a single job scheduler for all hosts on which jobs could be concurrently started. However, this approach greatly complicates the already complicated scheduling problem. Whereas the number of concurrent starts will be small compared to the total number of jobs, the semaphore approach was selected.

It is the intent of the developers that PBS will be expanded to encompass the concept of a single job whose execution is distributed among multiple hosts.

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 job 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.

The server calculates the *Resource Cost* of a job by summing the product of the amount of each resource multiplied by an assigned cost of the resource. A general system surcharge may also be assigned and added to the above sum. Resources with a “size” unit are converted to megabytes before the multiplication to keep the number from becoming too large. See the server attributes `resources_cost` and `system_cost`.

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 amount 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.

### 3.2.3. 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. 9

When a server queues a job with a dependency type of `syncwith`, `after`, `afterok`, `afternotok`, or `afterany` 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. 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. 9

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. 9

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. 9

The result is a pairing between corresponding `before...` and `after...` dependency types. 9

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. 9

### 3.3. Deferred Services 9

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. 9

The following rules apply to deferred services on behalf of jobs: 9

- If the server *cannot* complete a deferred service for a reason which is permanent, then the job is aborted. 9
- If the service cannot be completed at the current time but may be later, the service is retried a finite number of times. 9

#### 3.3.1. Job Scheduling 9

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: 9

- 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`. 9
- Termination of a running job. The termination may be normal execution completion, or because the job was deleted by request. 9
- Elapse of a specified cycle time as established by the administrator. 9
- 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. 9

The Scheduler is then treated as a privileged client and may make any request of the Server, including Run Job, Delete Job, Hold Job, or Modify Job/Queue/Server. 9

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. 9

### 3.3.2. File Staging 9

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: 9

```
-W stagein=local_file@host:remote_path[,local_file@host:remote_path,...]9
-W stageout=local_file@host:remote_path[,local_file@host:remote_path,...]9
```

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. 9

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. 9

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. 9

### 3.3.3. Job Initiation 9

Job initiation is to place a job into execution. The server creates a session leader that runs the shell program indicated by the Shell\_Path\_List attribute of the job. The path-name 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 server will search its execution path, \$PATH, for the shell. If the path name of the shell is omitted or is the null string, the server uses the login shell for the user under whose name the job is to be run. 9

The server will determine the user name under which the job is to be run by the following rules: 9

1. Select the user identifier from the User\_List job attribute which has a host name that matches the execution host. 9
2. Select the user identifier from the User\_List job attribute which has no associated host name. 9
3. Use the user name from the job\_owner attribute of the job. 9

The server will place the job into **running** state. 9

The server will create, in the environment of the session leader of the job, the environment variables named: 9

- PBS\_ENVIRONMENT** - the value of which is the string PBS\_BATCH. 9
- PBS\_QUEUE** - the value of which is the name of the execution queue. 9

The server 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. 9  
9

The server will place the required limits on the resources for which the host system supports resource limits. 9  
9

If the job had been run before and is now being *rerun*, the server will insure that the standard output and standard error streams of the job are appended to the prior streams, if any. 9  
9

If the server and host system support accounting, the server will use the value of the Account\_Name job attribute as required by the host system. 9  
9

If the server and host system support checkpoint, the server 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. 9  
9  
9  
9

The server will set up the standard output stream and the standard error stream of the job according to the following rules: 9  
9

- The stream will be located either (1) in a temporary file in the server's spool directory, or (2) a file in the user's home directory. The choice is determined by a server build time configuration parameter. 9  
9
- If the job attribute Join\_Path has the value eo or the value oe, the server connects the standard error stream of the job to the same file as the standard output stream. 9  
9

If the value of the job attribute Mail\_Points contains the value {beginning}, the server will send mail to each mail address specified in the job attribute Mail\_Users. 9  
9

### 3.3.4. Job Routing 9

Job routing is moving a job from a routing queue to one of the destinations associated with the queue. 9  
9

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. 9  
9  
9

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. 9  
9

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. 9  
9

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 number of tries in the queue attribute number\_retries. 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. 9  
9  
9  
9

### 3.3.5. Job Exit 9

When the session leader of a batch job exits, the server will perform the following actions in the order listed. 9  
9

Place the job in the **exiting** state. 9

“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. 9

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 steam of the job to the home directory of the user under whose name the job executed. The file name for the output is 9

job\_name.oseq\_number 9

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 9

job\_name.eseq\_number 9

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. 9

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. 9

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. 9

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. 9

If out staging of files is supported, the files listed in the outfile resource will be copied to the specified destination. 9

The job will be removed from the execution queue. 9

**3.3.6. Job Aborts** 9

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. 9

The job is removed from the queue. 9

**3.3.7. Timed Events** 9

The server performs certain events at a specified time or after a specified time delay. 9

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. 9

If the server is unable to make connection with another server, it is to retry after a time specified either by the routing queue attribute route\_retry\_time, or the general server attribute network\_retry\_time. 9

**3.3.8. Event Logging** 9

The various daemons including the **PBS** server will maintain a log file of events. This file is available to the batch administrator for analysis of past events. 9



The file will be maintained under the path name {PBS\_SERVER\_HOME}/server\_log/date, where date is the date in the form yyymmdd when the log file started (see the -L option in pbs\_server(8B)).

The events recorded by the server in the file are specified by the server attribute log\_events which is a bit string with each bit determining if a type of event is logged:

- 1 Internal PBS errors.
- 2 System (OS) errors such as malloc failed.
- 4 Administrator related events, such as changing server or queue attributes.
- 8 Job related events: submitted, ran, deleted, ...
- 16 (0x010) Job resource usage, this duplicates the accounting information in the log.
- 32 (0x020) Security related events, such as attempts to connect from an unknown host.
- 64 (0x040) When the scheduler was called and why.
- 128 (0x080) First level, common, debug messages.
- 256 (0x100) Second level, more rare, debug messages.

The log file is a text file with each entry terminated by a new line. The format of an entry is:

```
date          time;event_code;server_name;object_type;object_name;mes-
sage_text
```

The date time field is a date and time stamp in the format: mm/dd/yyyy hh:mm:ss. The event\_code is the type of event which triggered the event logging. It corresponding to the bit position, 0 to n, in the log\_events server attribute. The server\_name is the name of the server which logged the message. This is recorded in case a site wishes to merge and sort the various logs in a single file. The object\_type is the type of object which the message is about, Svr for server, Que for queue, Job for job, Req for request, or Fil for file. The object\_name is the name of the specific object. message\_text field is the text of the log message.

### 3.3.9. Accounting

The PBS server maintains an accounting file. The file will be maintained under the path name {PBS\_SERVER\_HOME}/server\_priv/accounting/day. Where day is the date in the form yyymmdd when the accounting file started (see the -A option in pbs\_server(8B)).

The account file is a text file with each entry terminated by a new line. The format of an entry is:

```
date time;record_type;id_string;message_text
```

The date time field is a date and time stamp in the format: mm/dd/yyyy hh:mm:ss. The id\_string is the job, reservation, or reservation-job identifier. The message\_text is ascii text. The content depends on the record type. The message text format is blank separated keyword=value fields. The record\_type is a single character indicating the type of record. The types are:

|   |   |   |
|---|---|---|
| A | Job was aborted by the server.  | 9 |
| D | Job was deleted by request. The message_text will contain requestor=us-<br>er@host to identify who deleted the job.   | 9 |
| E | Job ended (terminated execution). The message_text field contains:<br>user=username - the user name under which the job executed.<br>group=groupname - the group name under which the job executed.<br>account=account_string - if job has an "account name" string.<br>jobname=job_name - the name of the job.<br>queue=queue_name - the name of the queue from which the job is executed.<br>resvname=reservation_name - if job draws its resources from a resources<br>reservation and that reservation has a name.<br>resvID=reservation_ID_string - if job draws its resources from a resources<br>reservation.<br>resvjobID=reservation_ID_string - if job is as a "reservation-job" (advance<br>reservation of resources).<br>ctime=time - time in seconds when job was created (first submitted).<br>qtime=time - time in seconds when job was queued into current queue.<br>etime=time - time in seconds when job became eligible to run; no holds, etc.<br>start=time - time in seconds when job execution started.<br>exec_host=host - name of host on which the job is being executed.<br>Resource_List.resource=limit - list of the specified resource limits.<br>session=sesid - session number of job.<br>alt_id=id - Optional alternate job identifier. Will be included only for cer-<br>tain systems:<br>Irix 6.x with Array Services – The alternate id is the Array Session<br>Handle (ASH) assigned to the job.<br>end=time - time in secnds when job ended execution.<br>Exit_status=value - the exit status of the job. If the value is less than<br>10000 (decimal) it is the exit value of the top level process of the job, typically<br>the shell. If the value is greater than 10000, the top process exited on a signal<br>whose number is given by subtracting 10000 from the exit value.<br>Resources_used.resource=limit - list of the specified resource limits.<br>For Resource_List and Resources_used, there is one entry per resource. | 9 |
| C | Job was checkpointed and held.  | 9 |
| Q | Job entered a queue. The message_text contains queue=name identifying the<br>queue into which the job was placed. There will be a new Q record each time<br>the job is routed or moved to a new (or the same) queue.  | 9 |
| R | Job was rerun.  | 9 |
| S | Job execution started. The message_text field contains:<br>user=username - the user name under which the job executed.<br>group=groupname - the group name under which the job executed.<br>jobname=job_name - the name of the job.<br>queue=queue_name - the name of the queue from which the job is executed.<br>ctime=time - time in seconds when job was created (first submitted).<br>qtime=time - time in seconds when job was queued into current queue.<br>etime=time - time in seconds when job became eligible to run; no holds, etc.<br>start=time - time in seconds when job execution started.<br>exec_host=host - name of host on which the job is being executed.<br>Resource_List.resource=limit - list of the specified resource limits.<br>session=sesid - session number of job.   | 9 |
| T | Job was restarted from a checkpoint file.   | 9 |

- K Scheduler or server requested removal of the reservation. The message\_text field contains: requestor=user@host - to identify who deleted the resources reservation. 9
- k Resources reservation terminated by ordinary client - e.g. an owner issuing a pbs\_rdel command. The message\_text field contains: requestor=user@host - to identify who deleted the resources reservation. 9
- U Created unconfirmed resources reservation on server. The message\_text field contains either: 9  
     requestor=user@host - to identify who requested the resources reservation. 9  
     or 9  
     requestor=user@host - to identify who requested the resources reservation, the specification of how long he is willing to "be on hold" for confirmation and, the handling of the request should confirmation not occur in the time allotted - see pbs\_rsub(7B) option [-I]. 9
- Y Resources reservation confirmed by the scheduler. The message\_text field contains the same item (items) as in a U record type. 9
- F Resources reservation period finished. 9
- B Beginning of reservation's period. The message\_text field contains: 9  
     owner=ownername - name of party who submitted the resources reservation request. 9  
     name=reservation\_name - if submitter supplied a name string for the reservation. 9  
     account=account\_string - if submitter supplied a to be recorded in accounting. 9  
     queue=queue\_name - the name of the instantiated reservation queue if this is a general resources reservation. If the resources reservation is for a reservation job, this is the name of the queue to which the reservation-job belongs. 9  
     ctime=creation\_time - time at which the resources reservation got created, seconds since the epoch. 9  
     start=period\_start - time at which the reservation period is to start, in seconds since the epoch. 9  
     end=period\_end - time at which the reservation period is to end, seconds since the epoch. 9  
     duration=reservation\_duration - the duration specified or computed for the resources reservation, in seconds. 9  
     nodes=reserved\_nodes - if nodes with specified properties are required, this string is the allocated set. 9  
     authorized\_users=users\_list - the list of acl\_users on the queue that is instantiated to service the reservation. 9  
     authorized\_groups=groups\_list - if specified, the list of acl\_groups on the queue that is instantiated to service the reservation. 9  
     authorized\_hosts=hosts\_list - if specified, the list of acl\_hosts on the queue that is instantiated to service the reservation. 9  
     resource\_list=resources\_list - list of resources requested by the reservation. Resources are listed individually as, for example: resource\_list.ncpus=16 resource\_list.mem=1048676. 9

### 3.4. Resource Management 9

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. 9

3.4.1. Non Reservable Resources 9

Most Unix systems do not provide for resource reservation, only for limits. Resources, like memory, disk space, and cpu time, are handed out by the kernel on a first come first served basis. When a request exceeds the users limits, the request is denied or the job is signaled. To add resource reservation to a system is generally a major undertaking. One example is the Session Reservable File System, SRFS, extension to Unicos® developed at NAS. This extension required several additions to the kernel.

For resources which are not reservable, **PBS** manages resource allocation based on the amount "allocated" to **PBS** by the administrator. This available amount of each resource is maintained in the server attribute `resources_available`. The share of the resource "distributed" to each queue are maintained in an attribute for each of those objects. This attribute limits the aggregate total of the resources used by the jobs running under each object. This allocation is made by the batch system administrator. Most host systems do not provide support for dynamically adjusting the allocation to the batch system. A site may build in a procedure for adjusting the values of certain resources based on system load.

An example of non reservable resources is memory. The host operating system kernel manages memory, dynamically assigning physical memory to running processes. A limit can be set which the process cannot exceed, but memory cannot be reserved for a particular process in advance.

The `resources_max` attribute for the server and queue declares the maximum amount of the resource that a single job may be allocated. The server's `resources_max` is examined if there is not a `resources_max` value for the type of resource defined at the queue level.

**PBS** insures that the resources requested by a job fall with in the two groups of limits before the job is initiated.

3.4.2. Reservable Resources 9

On some hosts, certain resources types may be requested and the amount guaranteed to the process. Session Reservable File System, SRFS, is such a resource. For these types of resources, **PBS** will attempt to reserve the requested amount before scheduling the job for execution.

When the request to reserve resources is denied by the system, the job selection function may assist or *expedite* the job. In this case resources allocated to the job are not "released". **PBS** will attempt to acquire the remaining resources until it is successful and the job can be initiated, or until a time limit, specified by the queue attribute `reserved_expedite`, is reached. At that point all the resources are released. If a job is being expedited, other jobs whose resources do not conflict with the needs of the expedited job may be scheduled for execution.

When the reservable resources have been allocated to the job and the non reservable resources fit into what **PBS** has available, the job will be placed into execution.

3.4.3. Resource Limits 9

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: 9

1. The current queue's attribute resources\_default. 9
2. The server's attribute resources\_default. 9
3. The current queue's attribute resources\_max. 9
4. The server's attribute resources\_max. 9

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. 9

3.4.4. Resource Names 9

For additional information or information specific to an architecture, please see the man page for PBS Resources on that architecture. 9

The resources values take the following units: 9

**time** specifies a maximum time period the resource can be used. Time is expressed in seconds as an integer, or in the form: 9  
 [[hours:]minutes:]seconds[.milliseconds] 9  
 If specified, milliseconds are rounded to the nearest second. 9

**size** specifies the maximum amount in terms of bytes or words. It is expressed in the form integer[suffix] The suffix is a multiplier defined in the following table, "b" means bytes (the default) and "w" means words. The size of a word is calculated on the execution server as its word size.

| Suffix |    | Multiplier        |
|--------|----|-------------------|
| b      | w  | 1                 |
| kb     | kw | 1024              |
| mb     | mw | 1,048,576         |
| gb     | gw | 1,073,741,824     |
| tb     | tw | 1,099,511,627,776 |

**string** of characters which must be interpreted by the execution server. It is frequently a path name.

**unitary** The maximum amount of a resource which is expressed as a simple integer. 9

**4. Interface Library** 9


Part of the **PBS** package 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. 9

Users may have requirements to build jobs which could status itself or spawn off new jobs. Or they may wish to customize the job status display rather than use **qs-tat**. Administrators may use the interface library to build new control commands. 9

**4.1. Interface Library Overview** 9

The IFL provides a user callable function which corresponds 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`. 9

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 *Interface Facility*, IFF, as shown in figure 4-1. The purpose of the IFF is to provide the user a credential which validates the user's identity. This credential is included in each batch request. The IFF provided credential prevents a user from spoofing another user's identity. See chapter 10, Security, for more detail. 9



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

After all requests have been made to a server, its connection is closed via a call to `pbs_disconnect()`. 9

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 insure the parameters are in the correct syntax. 9

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`. 9

The library routine will accept the parameters and build the corresponding batch request. This request is then passed to the server communication routine. 9

**4.2. Interface Library Routines** 9

The following "man" pages describe the user callable functions in the IFL. The func- 9  
tions are found in files src/lib/Libifl/pbsD\_\*.c. 9

|   |   |
|---|---|
| 4.2.1. Alter Job  | 9 |
| <b>NAME</b>   | 9 |
| pbs_alterjob - alter pbs batch job  | 9 |
| <b>SYNOPSIS</b>   | 9 |
| #include <pbs_error.h>  | 9 |
| #include <pbs_ifl.h>  | 9 |
| <b>int pbs_alterjob(int connect, char *job_id, struct attrl *attrib, char *extend)</b>  | 9 |
|   | 9 |
| <b>DESCRIPTION</b>  | 9 |
| Issue a batch request to alter a batch job.   | 9 |
| A <i>Modify Job</i> batch request is generated and sent to the server over the connection specified by <i>connect</i> which is the return value of <b>pbs_connect()</b> .                     | 9 |
| The argument, <i>job_id</i> , identifies which job is to be altered, it is specified in the form:   | 9 |
| sequence_number.server  | 9 |
| The parameter, <i>attrib</i> , is a pointer to an <i>attrl</i> structure which is defined in pbs_ifl.h as:  | 9 |
|   | 9 |
| struct attrl {  | 9 |
| char *name;   | 9 |
| char *resource;   | 9 |
| char *value   | 9 |
| struct attrl *next;   | 9 |
| };  | 9 |
| The <i>attrib</i> list is terminated by the first entry where <i>next</i> is a null pointer.  | 9 |
| 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: | 9 |
| #define ATTR_a "Execution_Time"   | 9 |
| Alter the job's execution time.   | 9 |
| #define ATTR_A "Account_Name"   | 9 |
| Alter the account string.   | 9 |
| #define ATTR_c "Checkpoint"   | 9 |
| Alter the checkpoint interval.  | 9 |
| #define ATTR_e "Error_Path"   | 9 |
| Alter the path name for the standard error of the job.  | 9 |
| #define ATTR_g "Group_List"   | 9 |
| Alter the list of group names under which the job may execute.  | 9 |
| #define ATTR_h "Hold_Types"   | 9 |
| Alter the hold types.   | 9 |
| #define ATTR_j "Join_Path"  | 9 |
| Alter if standard error and standard output are joined (merged).  | 9 |
| #define ATTR_k "Keep_Files"   | 9 |
| Alter which output of the job is kept on the execution host.  | 9 |
| #define ATTR_l "Resource_List"  | 9 |
| Alter the value of a named resource.  | 9 |



|  |   |
|--|---|
| #define ATTR_m "Mail_Points"   | 9 |
| Alter the points at which the server will send mail about the job.   | 9 |
| #define ATTR_M "Mail_Users"  | 9 |
| Alter the list of users who would receive mail about the job.  | 9 |
| #define ATTR_N "Job_Name"  | 9 |
| Alter the job name.  | 9 |
| #define ATTR_o "Output_Path"   | 9 |
| Alter the path name for the standard output of the job.  | 9 |
| #define ATTR_p "Priority"  | 9 |
| Alter the priority of the job.   | 9 |
| #define ATTR_r "Rerunable"   | 9 |
| Alter the rerunable flag.  | 9 |
| #define ATTR_S "Shell_Path_List"   | 9 |
| Alter the path to the shell which will interprets the job script.  | 9 |
| #define ATTR_u "User_List"   | 9 |
| Alter the list of user names under which the job may execute.  | 9 |
| #define ATTR_v "Variable_List"   | 9 |
| Alter the list of environmental variables which are to be exported to the job.   | 9 |
| #define ATTR_depend "depend"   | 9 |
| Alter the inter-job dependencies.  | 9 |
| #define ATTR_stagein "stagein"   | 9 |
| Alter the list of files to be staged-in before job execution.  | 9 |
| #define ATTR_stageout "stageout"   | 9 |
| Alter the list of files to be staged-out after job execution.  | 9 |
| If <i>attrib</i> itself is a null pointer, then no attributes are altered.   | 9 |
| Associated with an attribute of type ATTR_l (the letter ell) is a resource name indicated by <i>resource</i> in the <i>attrl</i> structure. All other attribute types should have a pointer to a null string ("") for <i>resource</i> .      | 9 |
| 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.                             | 9 |
| Certain attributes of a job may or may not be alterable depending on the state of the job; see <b>qalter</b> (1B).   | 9 |
| The parameter, <i>extend</i> , is reserved for implementation defined extensions. It is not currently used by this function.   | 9 |
| <b>SEE ALSO</b>  | 9 |
| qalter(1B), qhold(1B), qrls(1B), qsub(1B), pbs_connect(3B), pbs_holdjob(3B), and pbs_rl-sjob(3B)   | 9 |
| <b>DIAGNOSTICS</b>   | 9 |
| When the batch request generated by <b>pbs_alterjob</b> () 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. | 9 |

|   |   |
|---|---|
| 4.2.2. Connect with Server  | 9 |
| <b>NAME</b>   | 9 |
| <code>pbs_connect</code> - connect to a pbs batch server  | 9 |
| <b>SYNOPSIS</b>   | 9 |
| <code>#include &lt;pbs_error.h&gt;</code>   | 9 |
| <code>#include &lt;pbs_ifl.h&gt;</code>   | 9 |
| <b>int pbs_connect(char *server)</b>  | 9 |
| <b>extern char *pbs_server;</b>   | 9 |
| <b>DESCRIPTION</b>  | 9 |
| A virtual stream (TCP/IP) connection is established with the server specified by <i>server</i> .  | 9 |
| This function must be called before any of the other <b>pbs_</b> 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.  | 9 |
| The connection should be closed by a call to <b>pbs_disconnect()</b> when all requests have been sent to the server.  | 9 |
| The parameter, <i>server</i> , is of the form <code>host_name[:port]</code> , see section 2.7.9. If <code>port</code> is not specified, the standard PBS port number will be used.  | 9 |
| If the parameter, <i>server</i> , is either the null string or a null pointer, a connection will be opened to the default server. The default server is defined in section 2.7.4.   | 9 |
| The variable <i>pbs_server</i> , declared in <code>pbs_ifl.h</code> , is set on return to point to the server name to which <code>pbs_connect()</code> connected or attempted to connect.   | 9 |
| <b>SEE ALSO</b>   | 9 |
| <code>qsub(1B)</code> , <code>pbs_alterjob(3B)</code> , <code>pbs_deljob(3B)</code> , <code>pbs_disconnect(3B)</code> , <code>pbs_geterrmsg(3B)</code> , <code>pbs_holdjob(3B)</code> , <code>pbs_locate(3B)</code> , <code>pbs_manager(3B)</code> , <code>pbs_movejob(3B)</code> , <code>pbs_msgjob(3B)</code> , <code>pbs_rerunjob(3B)</code> , <code>pbs_rlsjob(3B)</code> , <code>pbs_runjob(3B)</code> , <code>pbs_selectjob(3B)</code> , <code>pbs_selstat(3B)</code> , <code>pbs_sigjob(3B)</code> , <code>pbs_statjob(3B)</code> , <code>pbs_statque(3B)</code> , <code>pbs_statserver(3B)</code> , <code>pbs_submit(3B)</code> , <code>pbs_terminate(3B)</code> , <code>pbs_server(8B)</code> , and the PBS External Reference Specification | 9 |
| <b>DIAGNOSTICS</b>  | 9 |
| 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 <code>pbs_errno</code> .  | 9 |

|  |   |
|--|---|
| 4.2.3. Delete Job  | 9 |
| <b>NAME</b>  | 9 |
| pbs_deljob - delete a pbs batch job  | 9 |
| <b>SYNOPSIS</b>  | 9 |
| #include <pbs_error.h>   | 9 |
| #include <pbs_ifl.h>   | 9 |
| <b>int pbs_deljob(int connect, char *job_id, char *extend)</b>   | 9 |
| <b>DESCRIPTION</b>   | 9 |
| Issue a batch request to delete a batch job. If the batch job is running, the execution server will send the <b>SIGTERM</b> signal followed by <b>SIGKILL</b> .  | 9 |
| A <i>Delete Job</i> batch request is generated and sent to the server over the connection specified by <i>connect</i> which is the return value of <b>pbs_connect()</b> .  | 9 |
| The argument, <i>job_id</i> , identifies which job is to be deleted, it is specified in the form: <i>sequence_number.server</i>  | 9 |
| The argument, <i>extend</i> , is overloaded to serve two purposes. If the pointer <i>extend</i> points to a string of the form:<br>deldelay= <i>nnnn</i> ,   | 9 |
| it is used to provide control over the delay between sending <b>SIGTERM</b> and <b>SIGKILL</b> signals to a running job. The characters <i>nnnn</i> specify a unsigned decimal integer time delay in seconds. If <i>extend</i> is the null pointer or points to a null string, the administrator established default time delay is used. | 9 |
| If <i>extend</i> 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.   | 9 |
| <b>SEE ALSO</b>  | 9 |
| qdel(1B) and pbs_connect(3B)   | 9 |
| <b>DIAGNOSTICS</b>   | 9 |
| When the batch request generated by the <b>pbs_deljob()</b> 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 <i>pbs_errno</i> .  | 9 |

|   |   |
|---|---|
| 4.2.4. Disconnect from Server   | 9 |
| <b>NAME</b>   | 9 |
| <code>pbs_disconnect</code> - disconnect from a pbs batch server  | 9 |
| <b>SYNOPSIS</b>   | 9 |
| <code>#include &lt;pbs_error.h&gt;</code>   | 9 |
| <code>#include &lt;pbs_ifl.h&gt;</code>   | 9 |
| <b>int pbs_disconnect(int connect)</b>  | 9 |
| <b>DESCRIPTION</b>  | 9 |
| The virtual stream connection specified by <i>connect</i> , which was established with a server by a call to <b>pbs_connect()</b> , is closed.  | 9 |
| <b>SEE ALSO</b>   | 9 |
| <code>pbs_connect(3B)</code>  | 9 |
| <b>DIAGNOSTICS</b>  | 9 |
| 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 <code>pbs_errno</code> . | 9 |

|   |   |
|---|---|
| 4.2.5. Get Error Message Text   | 9 |
| <b>NAME</b>   | 9 |
| <code>pbs_geterrmsg</code> - get error message for last pbs batch operation   | 9 |
| <b>SYNOPSIS</b>   | 9 |
| <code>#include &lt;pbs_error.h&gt;</code>   | 9 |
| <code>#include &lt;pbs_ifl.h&gt;</code>   | 9 |
| <b>char *pbs_geterrmsg(int connect)</b>   | 9 |
| <b>DESCRIPTION</b>  | 9 |
| Return the error message text associated with a batch server request.   | 9 |
| Because different server implementations may have different error conditions, the batch request - reply protocol allows for the server to include the text of an error message in the reply to a batch request.   | 9 |
| If the preceding batch interface library call over the connection specified by <i>connect</i> 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. | 9 |
| <b>SEE ALSO</b>   | 9 |
| <code>pbs_connect(3B)</code>  | 9 |
| <b>DIAGNOSTICS</b>  | 9 |
| If an error text message was returned by a server in reply to the previous call to a batch interface library function, <b>pbs_geterrmsg()</b> will return a pointer to it. Otherwise, <code>pbs_geterrmsg()</code> returns the null pointer.                              | 9 |

|  |   |
|--|---|
| 4.2.6. Hold Job  | 9 |
| <b>NAME</b>  | 9 |
| <code>pbs_holdjob</code> - place a hold on a pbs batch job   | 9 |
| <b>SYNOPSIS</b>  | 9 |
| <code>#include &lt;pbs_error.h&gt;</code>  | 9 |
| <code>#include &lt;pbs_ifl.h&gt;</code>  | 9 |
| <b>int pbs_holdjob(int connect, char *job_id, char *hold_type, char *extend)</b>   | 9 |
| <b>DESCRIPTION</b>   | 9 |
| Issue a batch request to place a hold upon a job.  | 9 |
| A <i>Hold Job</i> batch request is generated and sent to the server over the connection specified by <i>connect</i> which is the return value of <b>pbs_connect()</b> .  | 9 |
| The argument, <i>job_id</i> , identifies which job is to be held, it is specified in the form: <i>sequence_number.server</i>   | 9 |
| The parameter, <i>hold_type</i> , contains the type of hold to be applied. The possible values are defined in <code>pbs_ifl.h</code> as:   | 9 |
| <code>#define USER_HOLD "u"</code>   | 9 |
| Available to the owner of the job, the batch operator, and the batch administrator.  | 9 |
| <code>#define OTHER_HOLD "o"</code>  | 9 |
| Available to the batch operator and the batch administrator.   | 9 |
| <code>#define SYSTEM_HOLD "s"</code>   | 9 |
| Available only to the batch administrator.   | 9 |
| If <i>hold_type</i> is either a null pointer or points to a null string, <code>USER_HOLD</code> will be applied.   | 9 |
| The parameter, <i>extend</i> , is reserved for implementation defined extensions. It is not currently used by this function.   | 9 |
| <b>SEE ALSO</b>  | 9 |
| <code>qhold(1B)</code> , <code>pbs_connect(3B)</code> , <code>pbs_alterjob(3B)</code> , and <code>pbs_rlsjob(3B)</code>  | 9 |
| <b>DIAGNOSTICS</b>   | 9 |
| When the batch request generated by <b>pbs_holdjob()</b> 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 <code>pbs_errno</code> . | 9 |

|  |   |
|--|---|
| 4.2.7. Locate Job  | 9 |
| <b>NAME</b>  | 9 |
| <code>pbs_locjob</code> - locate current location of a pbs batch job   | 9 |
| <b>SYNOPSIS</b>  | 9 |
| <code>#include &lt;pbs_error.h&gt;</code>  | 9 |
| <code>#include &lt;pbs_ifl.h&gt;</code>  | 9 |
| <b><code>char *pbs_locjob(int connect, char *job_id, char *extend)</code></b>  | 9 |
| <b>DESCRIPTION</b>   | 9 |
| 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.   | 9 |
| A <i>Locate Job</i> batch request is generated and sent to the server over the connection specified by <i>connect</i> which is the return value of <b><code>pbs_connect()</code></b> .   | 9 |
| The argument, <i>job_id</i> , identifies which job is to be located, it is specified in the form: <code>sequence_number.server</code>  | 9 |
| The argument, <i>extend</i> , is reserved for implementation defined extensions. It is not currently used by this function.  | 9 |
| The return value is a pointer to a character sting which contains the current location if known. The syntax of the location string is: <code>queue@server_name</code> . If the location of the job is not known, the return value is the NULL pointer.   | 9 |
| <b>SEE ALSO</b>  | 9 |
| <code>qsub(1B)</code> and <code>pbs_connect(3B)</code>   | 9 |
| <b>DIAGNOSTICS</b>   | 9 |
| When the batch request generated by the <b><code>pbs_locjob()</code></b> 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 <code>pbs_errno</code> . | 9 |

## 4.2.8. Manager Server 9

**NAME** 9

pbs\_manager - administrator a pbs batch object 9

**SYNOPSIS** 9

#include &lt;pbs\_error.h&gt; 9

#include &lt;pbs\_ifl.h&gt; 9

**int pbs\_manager(int connect, int command, int obj\_type, char \*obj\_name,  
struct attrop1 \*attrib, char \*extend)** 9**DESCRIPTION** 9

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: 9

#define MGR\_CMD\_CREATE 0 9

#define MGR\_CMD\_DELETE 1 9

#define MGR\_CMD\_SET 2 9

#define MGR\_CMD\_UNSET 3 9

The parameter, *obj\_type*, declares the type of object upon which the command operates, see pbs\_ifl.h: 9

#define MGR\_OBJ\_SERVER 0 9

#define MGR\_OBJ\_QUEUE 1 9

The parameter, *obj\_name*, is the name of the specific object. 9

The parameter, *attrib*, is a pointer to an *attrop1* structure which are defined in pbs\_ifl.h as: 9

```

struct attrop1 {
    char    *name;
    char    *resource;
    char    *value;
    enum batch_op op;
    struct attrop1 *next;
};

```

The *attrib* list is terminated by the first entry where *next* is a null pointer. 9

The *name* member points to a string which is the name of the attribute. 9

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. 9

The *value* member points to a string which is the new value of the attribute. 9

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 }; 9

The full range of *batch\_op* values is { SET, UNSET, INCR, DECR, EQ, NE, GE, GT, LE, LT } Only the SET, UNSET, INCR, and DECR are allowed in a manager call, others will be rejected by the server. 9



The parameter *extend* is reserved for implementation defined extensions. It is not currently used by this function. 9  
9

Functions MGR\_CMD\_CREATE and MGR\_CMD\_DELETE require PBS Manager privilege. 9  
Functions MGR\_CMD\_SET and MGR\_CMD\_UNSET require PBS Manager or Operator privilege. 9  
9

**SEE ALSO** 9  
qmgr(8B) and pbs\_connect(3B) 9

**DIAGNOSTICS** 9  
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**. 9  
9

|   |   |
|---|---|
| 4.2.9. Move Job   | 9 |
| <b>NAME</b>   | 9 |
| <code>pbs_movejob</code> - move a pbs batch job to a new destination  | 9 |
| <b>SYNOPSIS</b>   | 9 |
| <code>#include &lt;pbs_error.h&gt;</code>   | 9 |
| <code>#include &lt;pbs_ifl.h&gt;</code>   | 9 |
| <b>int pbs_movejob(int connect, char *job_id, char *destination, char *extend)</b>  | 9 |
| <b>DESCRIPTION</b>  | 9 |
| 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.  | 9 |
| A <i>Move Job</i> batch request is generated and sent to the server over the connection specified by <i>connect</i> which is the return value of <code>pbs_connect()</code> .   | 9 |
| The <i>job_id</i> parameter identifies which job is to be moved; it is specified in the form: <code>sequence_number.server</code>   | 9 |
| The <i>destination</i> parameter specifies the new destination for the job. It is specified as: <code>[queue][@server]</code> . If <i>destination</i> is a null pointer or a null string, the destination will be the default queue at the current server. If <i>destination</i> specifies a queue but not a server, the destination will be the named queue at the current server. If <i>destination</i> specifies a server but not a queue, the destination will be the default queue at the named server. If <i>destination</i> specifies both a queue and a server, the destination is that queue at that server. | 9 |
| A job in the <b>Running</b> , <b>Transiting</b> , or <b>Exiting</b> state cannot be moved.  | 9 |
| The parameter, <i>extend</i> , is reserved for implementation defined extensions. It is not currently used by this function.  | 9 |
| <b>SEE ALSO</b>   | 9 |
| <code>qmove(1B)</code> , <code>qsub(1B)</code> , and <code>pbs_connect(3B)</code>   | 9 |
| <b>DIAGNOSTICS</b>  | 9 |
| When the batch request generated by <code>pbs_movejob()</code> 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 <code>pbs_errno</code> .  | 9 |

|   |        |
|---|--------|
| 4.2.10. Message Job   | 9      |
| <b>NAME</b>   | 9      |
| pbs_msgjob - record a message for a running pbs batch job   | 9      |
| <b>SYNOPSIS</b>   | 9      |
| #include <pbs_error.h>  | 9      |
| #include <pbs_ifl.h>  | 9      |
| <b>int pbs_msgjob(int connect, char *job_id, int file, char *message,<br/>char *extend)</b>   | 9<br>9 |
| <b>DESCRIPTION</b>  | 9      |
| Issue a batch request to write a message in an output file of a batch job.  | 9      |
| A <i>Message Job</i> batch request is generated and sent to the server over the connection specified by <i>connect</i> which is the return value of <b>pbs_connect()</b> .  | 9<br>9 |
| The argument, <i>job_id</i> , identifies the job to which the message is to be sent; it is specified in the form: <i>sequence_number.server</i>   | 9<br>9 |
| The parameter, <i>file</i> , indicates the file or files to which the message string is to be written. The following values are defined in pbs_ifl.h:   | 9<br>9 |
| #define MSG_ERR 2   | 9      |
| directs the message to the standard error stream of the job.  | 9      |
| #define MSG_OUT 1   | 9      |
| directs the message to the standard output stream of the job.   | 9      |
| The parameter, <i>message</i> , is the message string to be written.  | 9      |
| The parameter, <i>extend</i> , is reserved for implementation defined extensions. It is not currently used by this function.  | 9<br>9 |
| <b>SEE ALSO</b>   | 9      |
| qmsg(1B) and pbs_connect(3B)  | 9      |
| <b>DIAGNOSTICS</b>  | 9      |
| When the batch request generated by <b>pbs_msgjob()</b> 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. | 9<br>9 |

|   |   |
|---|---|
| 4.2.11. Order Job   | 9 |
| <b>NAME</b>   | 9 |
| pbs_orderjob - reorder pbs batch jobs in a queue  | 9 |
| <b>SYNOPSIS</b>   | 9 |
| #include <pbs_error.h>  | 9 |
| #include <pbs_ifl.h>  | 9 |
| <b>int pbs_orderjob(int connect, char *job_id1, char *job_id2, char *extend)</b>  | 9 |
| <b>DESCRIPTION</b>  | 9 |
| Issue a batch request to swap the order of two jobs with a single queue.  | 9 |
| An <i>Order Job</i> batch request is generated and sent to the server over the connection specified by <i>connect</i> which is the return value of <b>pbs_connect()</b> .   | 9 |
| The parameters <i>job_id1</i> and <i>job_id2</i> identify which jobs are to be swapped. They are specified in the form: <i>sequence_number.server</i> .   | 9 |
| The parameter, <i>extend</i> , is reserved for implementation defined extensions. It is not currently used by this function.  | 9 |
| <b>SEE ALSO</b>   | 9 |
| qorder(1B), qmove(1B), qsub(1M), and pbs_connect(3B)  | 9 |
| <b>DIAGNOSTICS</b>  | 9 |
| When the batch request generated by <b>pbs_orderjob()</b> 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 <i>pbs_errno</i> . | 9 |

|   |   |
|---|---|
| 4.2.12. Rerun Job   | 9 |
| <b>NAME</b>   | 9 |
| pbs_rerunjob - rerun a pbs batch job  | 9 |
| <b>SYNOPSIS</b>   | 9 |
| #include <pbs_error.h>  | 9 |
| #include <pbs_ifl.h>  | 9 |
| <b>int pbs_rerunjob(int connect, char *job_id, char *extend)</b>  | 9 |
| <b>DESCRIPTION</b>  | 9 |
| Issue a batch request to rerun a batch job.   | 9 |
| A <i>Rerun Job</i> batch request is generated and sent to the server over the connection specified by <i>connect</i> which is the return value of <b>pbs_connect()</b> .  | 9 |
| If the job is marked as being not rerunable, the request will fail and an error will be returned.   | 9 |
| The argument, <i>job_id</i> , identifies which job is to be rerun it is specified in the form: <i>sequence_number.server</i>  | 9 |
| The parameter, <i>extend</i> , is reserved for implementation defined extensions. It is not currently used by this function.  | 9 |
| <b>SEE ALSO</b>   | 9 |
| qrerun(1B), qsub(1B), and pbs_connect(3B)   | 9 |
| <b>DIAGNOSTICS</b>  | 9 |
| When the batch request generated by <b>pbs_rerunjob()</b> 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 <i>pbs_errno</i> . | 9 |

|   |   |
|---|---|
| 4.2.13. Release Job   | 9 |
| <b>NAME</b>   | 9 |
| pbs_rlsjob - release a hold on a pbs batch job  | 9 |
| <b>SYNOPSIS</b>   | 9 |
| #include <pbs_error.h>  | 9 |
| #include <pbs_ifl.h>  | 9 |
| <b>int pbs_rlsjob(int connect, char *job_id, char *hold_type, char *extend)</b>   | 9 |
| <b>DESCRIPTION</b>  | 9 |
| Issue a batch request to release a hold from a job.   | 9 |
| A <i>Release Job</i> batch request is generated and sent to the server over the connection specified by <i>connect</i> which is the return value of <b>pbs_connect()</b> .  | 9 |
| The argument, <i>job_id</i> , identifies the job from which the hold is to be released, it is specified in the form: <i>sequence_number.server</i>  | 9 |
| The parameter, <i>hold_type</i> , contains the type of hold to be released. The possible values are defined in <i>pbs_ifl.h</i> as:   | 9 |
| #define USER_HOLD "u"   | 9 |
| Available to the owner of the job, the batch operator, and the batch administrator.   | 9 |
| #define OTHER_HOLD "o"  | 9 |
| Available to the batch operator and the batch administrator.  | 9 |
| #define SYSTEM_HOLD "s"   | 9 |
| Available only to the batch administrator.  | 9 |
| If <i>hold_type</i> is either a null pointer or points to a null string, USER_HOLD will be released.  | 9 |
| The parameter, <i>extend</i> , is reserved for implementation defined extensions. It is not currently used by this function.  | 9 |
| <b>SEE ALSO</b>   | 9 |
| qrls(1B), qhold(1B), qalter(1B), pbs_alterjob(3B), pbs_connect(3B), and pbs_holdjob(3B)   | 9 |
| <b>DIAGNOSTICS</b>  | 9 |
| When the batch request generated by <b>pbs_rlsjob()</b> 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 <i>pbs_errno</i> . | 9 |

|   |   |
|---|---|
| 4.2.14. Run Job   | 9 |
| <b>NAME</b>   | 9 |
| pbs_runjob - run a pbs batch job  | 9 |
| <b>SYNOPSIS</b>   | 9 |
| #include <pbs_error.h>  | 9 |
| #include <pbs_ifl.h>  | 9 |
| <b>int pbs_runjob(int connect, char *job_id, char *location, char *extend)</b>  | 9 |
| <b>int pbs_asyruntime(int connect, char *job_id, char *location, char *extend)</b>  | 9 |
| <b>DESCRIPTION</b>  | 9 |
| Issue a batch request to run a batch job.   | 9 |
| For <i>pbs_runjob()</i> a "Run Job" batch request is generated and sent to the server over the connection specified by <i>connect</i> which is the return value of <b>pbs_connect()</b> . 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. | 9 |
| For <i>pbs_asyruntime()</i> an "Asynchronous Run Job" request is generated and set 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.             | 9 |
| These requests requires that the issuing user have operator or administrator privilege.   | 9 |
| The argument, <i>job_id</i> , identifies which job is to be run it is specified in the form: <i>sequence_number.server</i>  | 9 |
| The argument, <i>location</i> , if not the null pointer or null string, specifies the location where the job should be run. The location is the name of a host in the the cluster managed by the server. If the server does not understand the location, it will reject the request and return an error.  | 9 |
| The argument, <i>extend</i> , is reserved for implementation defined extensions. It is not currently used by these functions.   | 9 |
| <b>SEE ALSO</b>   | 9 |
| qrun(8B), qsub(1B), and pbs_connect(3B)   | 9 |
| <b>DIAGNOSTICS</b>  | 9 |
| When the batch request generated by the <b>pbs_runjob()</b> or <b>pbs_asyruntime()</b> 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 <i>pbs_errno</i> .  | 9 |

## 4.2.15. Select Job 9

**NAME** 9

pbs\_selectjob - select pbs batch jobs 9

**SYNOPSIS** 9

#include &lt;pbs\_error.h&gt; 9

#include &lt;pbs\_ifl.h&gt; 9

**char \*\*pbs\_selectjob(int connect, struct attropl \*attrib, char \*extend)** 9**DESCRIPTION** 9

Issue a batch request to select jobs which meet certain criteria. **pbs\_selectjob()** returns a array of job identifiers which met the criteria. 9

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. 9

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()**. 9

The argument, *attrib*, is a pointer to an *attropl* structure which is defined in pbs\_ifl.h as: 9

```

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

```

9  
9  
9  
9  
9  
9  
9

The *attrib* list is terminated by the first entry where *next* is a null pointer. 9

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 defined in pbs\_ifl.h: 9

```

#define ATTR_a "Execution_Time"
    Select based upon the job's execution time.
#define ATTR_A "Account_Name"
    Select (E) based upon the account string.
#define ATTR_c "Checkpoint"
    Select based upon the checkpoint interval.
#define ATTR_e "Error_Path"
    Select (E) based upon the name of the standard error file.
#define ATTR_g "Group_List"
    Select (E) based upon the list of group names under which the job may execute.
#define ATTR_h "Hold_Types"
    Select (E) based upon the hold types.
#define ATTR_j "Join_Paths"
    Select (E) based upon the value of the join list.
#define ATTR_k "Keep_Files"
    Select (E) based upon the value of the keep files list.

```

9  
9  
9  
9  
9  
9  
9  
9  
9  
9  
9  
9



```

#define ATTR_l "Resource_List"          9
    Select based upon the value of the resource named in resource.  9
#define ATTR_m "Mail_Points"           9
    Select (E) based upon the setting of the mail points attribute.  9
#define ATTR_M "Mail_Users"            9
    Select (E) based upon the list of user names to which mail will be sent.  9
#define ATTR_N "Job_Name"              9
    Select (E) based upon the job name.  9
#define ATTR_o "Output_Path"           9
    Select (E) based upon the name of the standard output file.  9
#define ATTR_p "Rriority"              9
    Select based upon the priority of the job.  9
#define ATTR_q "destination"           9
    Select based upon the specified destination. Jobs selected are restricted to those 9
    residing in the named queue. If destination is the null string, the default queue at 9
    the server is assumed.  9
#define ATTR_r "Rerunable"             9
    Select (E) based upon the rerunable flag.  9
#define ATTR_session "session_id"      9
    Select based upon the session id assigned to running jobs.  9
#define ATTR_S "Shell_Path_List"       9
    Select (E) based upon the execution shell list.  9
#define ATTR_u "User_List"             9
    Select (E) based upon the owner of the jobs.  9
#define ATTR_v "Variable_List"         9
    Select (E) based upon the list of environment variables.  9
#define ATTR_ctime "ctime"             9
    Select based upon the creation time of the job.  9
#define ATTR_depend "depend"           9
    Select based upon the list of job dependencies.  9
#define ATTR_mtime "mtime"            9
    Select based upon the last modification time of the job.  9
#define ATTR_qtime "qtime"            9
    Select based upon the time of the job was placed into the current queue.  9
#define ATTR_qtype "queue_type"        9
    Select (E) base on the type of queue in which the job resides.  9
#define ATTR_stagein "stagein"         9
    Select based upon the list of files to be staged-in.  9
#define ATTR_stageout "stageout"       9
    Select based upon the list of files to be staged-out.  9
#define ATTR_state "job_state"         9
    Select based upon the state of the jobs. State is not a job attribute, but is included 9
    here to allow selection.  9
The op member defines the operator in the logical expression:  9
    value operator current_value  9
The logical expression must evaluate as true for the job to be selected. The permissible val- 9
ues of op are defined in pbs_ifl.h as: enum batch_op { ..., EQ, NE, GE, GT, LE, 9
LT, ... };. The attributes marked with (E) in the description above may only be selected 9

```

with the equal, EQ, or not equal, NE, operators. 9

The full range of `batch_op` values is { SET, UNSET, INCR, DECR, EQ, NE, GE, GT, 9  
LE, LT }. Only the relational operators are allowed in a select call, and others will be re- 9  
jected by the server. 9

If `attrib` itself is a null pointer, then no selection is done on the basis of attributes. 9

The return value is a pointer to a null terminated array of character pointers. Each charac- 9  
ter pointer in the array points to a character string which is a *job\_identifier* in the form: se- 9  
quence\_number.server@server 9

The array is allocated by `pbs_selectjob` via **malloc()**. When the array is no longer needed, 9  
the user is responsible for freeing it by a call to **free()**. The space for the array and the job 9  
identifier strings is malloc-ed by **pbs\_selectjob()** as one allocation. The array of pointers 9  
starts at the address returned by `malloc()`, so `free()` will work. 9

The parameter, *extend*, is reserved for implementation defined extensions. It is not currently 9  
used by this function. 9

**SEE ALSO** 9  
`qselect(1B)`, `pbs_alterjob(3B)`, and `pbs_connect(3B)` 9

**DIAGNOSTICS** 9

When the batch request generated by `pbs_selectjob()` function has been completed successful- 9  
ly by a batch server, the routine will return a pointer to the array of job identifiers. If no jobs 9  
met the criteria, the first pointer in the array will be the null pointer. 9

If an error occurred, a null pointer is returned and the error is available in the global integer 9  
`pbs_errno`. 9

## 4.2.16. Status of Selected Jobs 9

**NAME** 9

pbs\_selstat - obtain status of selected pbs batch jobs 9

**SYNOPSIS** 9

#include &lt;pbs\_error.h&gt; 9

#include &lt;pbs\_ifl.h&gt; 9

```
struct batch_status *pbs_selstat(int connect, struct attropl *sel_list,
struct attrl *rattrib) 9
```

```
void pbs_statfree(struct batch_status *psj) 9
```

**DESCRIPTION** 9

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. 9

This function is a combination of **pbs\_selectjobs()** and **pbs\_statjob()**. It is an extension to the POSIX Batch standard. 9

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. 9

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()**. 9

The parameter, *sel\_list*, is a pointer to an *attropl* structure which is defined in *pbs\_ifl.h* as: 9

```
struct attropl { 9
    struct attropl *next; 9
    char *name; 9
    char *resource; 9
    char *value; 9
    enum batch_op op; 9
}; 9
```

The *sel\_list* list is terminated by the first entry where *next* is a null pointer. 9

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 defined in *pbs\_ifl.h*: 9

```
#define ATTR_a "Execution_Time" 9
    Select based upon the job's execution time. 9
#define ATTR_A "Account_Name" 9
    Select (E) based upon the account string. 9
#define ATTR_c "Checkpoint" 9
    Select based upon the checkpoint interval. 9
#define ATTR_e "Error_Path" 9
    Select (E) based upon the name of the standard error file. 9
#define ATTR_g "Group_List" 9
    Select (E) based upon the list of group names under which the job may execute. 9
#define ATTR_h "Hold_Types" 9
    Select (E) based upon the hold types. 9
```

|  |   |
|--|---|
| #define ATTR_j "Join_Paths"  | 9 |
| Select (E) based upon the value of the join list.  | 9 |
| #define ATTR_k "Keep_Files"  | 9 |
| Select (E) based upon the value of the keep files list.  | 9 |
| #define ATTR_l "Resource_List"   | 9 |
| Select based upon the value of the resource named in <code>resource</code> .   | 9 |
| #define ATTR_m "Mail_Points"   | 9 |
| Select (E) based upon the setting of the mail points attribute.  | 9 |
| #define ATTR_M "Mail_Users"  | 9 |
| Select (E) based upon the list of user names to which mail will be sent.   | 9 |
| #define ATTR_N "Job_Name"  | 9 |
| Select (E) based upon the job name.  | 9 |
| #define ATTR_o "Output_Path"   | 9 |
| Select (E) based upon the name of the standard output file.  | 9 |
| #define ATTR_p "Priority"  | 9 |
| Select based upon the priority of the job.   | 9 |
| #define ATTR_q "destination"   | 9 |
| Select based upon the specified destination. Jobs selected are restricted to those residing in the named queue. If destination is the null string, the default queue at the server is assumed. | 9 |
| #define ATTR_r "Rerunable"   | 9 |
| Select (E) based upon the rerunable flag.  | 9 |
| #define ATTR_session "session_id"  | 9 |
| Select based upon the session id assigned to running jobs.   | 9 |
| #define ATTR_S "Shell_Path_List"   | 9 |
| Select (E) based upon the execution shell list.  | 9 |
| #define ATTR_u "User_List"   | 9 |
| Select (E) based upon the owner of the jobs.   | 9 |
| #define ATTR_v "Variable_List"   | 9 |
| Select (E) based upon the list of environment variables.   | 9 |
| #define ATTR_ctime "ctime"   | 9 |
| Select based upon the creation time of the job.  | 9 |
| #define ATTR_depend "depend"   | 9 |
| Select based upon the list of job dependencies.  | 9 |
| #define ATTR_mtime "mtime"   | 9 |
| Select based upon the last modification time of the job.   | 9 |
| #define ATTR_qtime "qtime"   | 9 |
| Select based upon the time of the job was placed into the current queue.   | 9 |
| #define ATTR_qtype "queue_type"  | 9 |
| Select (E) base on the type of queue in which the job resides.   | 9 |
| #define ATTR_stagein "stagein"   | 9 |
| Select based upon the list of files to be staged-in.   | 9 |
| #define ATTR_stageout "stageout"   | 9 |
| Select based upon the list of files to be staged-out.  | 9 |
| #define ATTR_state "job_state"   | 9 |
| Select based upon the state of the jobs. State is not a job attribute, but is included here to allow selection.  | 9 |

The `op` member defines the operator in the logical expression: 9  
     value operator current\_value 9

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. The full range of `batch_op` values is SET, UNSET, INCR, DECR, EQ, NE, GE, GT, LE, and LT, Only the relational operators are allowed in a `selstat` call, and others will be rejected by the server. 9

If `sel_list` itself is a null pointer, then no selection is done on the basis of attributes. 9

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 a 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. 9

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 9

```
struct batch_status { 9
    struct batch_status *next; 9
    char *name; 9
    struct attrl *attribs; 9
    char *text; 9
} 9
```

The entry, `attribs`, is a pointer to a list of `attrl` structures defined in `pbs_ifl.h` as: 9

```
struct attrl { 9
    struct attrl *next; 9
    char *name; 9
    char *resource; 9
    char *value; 9
}; 9
```

It is up the user to free the list of `batch_status` structures when no longer needed, by calling **`pbs_statfree()`**. 9

#### SEE ALSO 9

`qselect(1B)`, `pbs_alterjob(3B)`, `pbs_connect(3B)`, `pbs_statjob(3B)`, and `pbs_selectjob(3B)`. 9

#### DIAGNOSTICS 9

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_errno` will be set to a non-zero value. 9

|   |   |
|---|---|
| 4.2.17. Signal Job  | 9 |
| <b>NAME</b>   | 9 |
| pbs_sigjob - send a signal to a pbs batch job   | 9 |
| <b>SYNOPSIS</b>   | 9 |
| #include <pbs_error.h>  | 9 |
| #include <pbs_ifl.h>  | 9 |
| <b>int pbs_sigjob(int connect, char *job_id, char *signal, char *extend)</b>  | 9 |
| <b>DESCRIPTION</b>  | 9 |
| Issue a batch request to send a signal to a batch job.  | 9 |
| A <i>Signal Job</i> batch request is generated and sent to the server over the connection specified by <i>connect</i> which is the return value of <b>pbs_connect()</b> . If the batch job is in the <b>running</b> state, the batch server will send the job the signal number corresponding to the signal named in <i>signal</i> .  | 9 |
| The argument, <i>job_id</i> , identifies which job is to be signaled, it is specified in the form: <i>sequence_number.server</i>  | 9 |
| The <i>signal</i> 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, <b>suspend</b> and <b>resume</b> . 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 <b>running</b> state, no signal is sent and an error is returned, except when the signal is <b>resume</b> and the job is suspended. | 9 |
| The parameter, <i>extend</i> , is reserved for implementation defined extensions. It is not currently used by this function.  | 9 |
| <b>SEE ALSO</b>   | 9 |
| qsig(1B) and pbs_connect(3B)  | 9 |
| <b>DIAGNOSTICS</b>  | 9 |
| When the batch request generated by <b>pbs_sigjob()</b> 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.   | 9 |

|  |                  |
|--|------------------|
| 4.2.18. Stage In Files   | 9                |
| <b>NAME</b>  | 9                |
| pbs_stagein - request that files for a pbs batch job be staged in.   | 9                |
| <b>SYNOPSIS</b>  | 9                |
| #include <pbs_error.h>   | 9                |
| #include <pbs_ifl.h>   | 9                |
| <b>int pbs_stagein(int connect, char *job_id, char *location,<br/>char *extend)</b>  | 9<br>9           |
| <b>DESCRIPTION</b>   | 9                |
| Issue a batch request to start the stage in of files specified in the stagein attribute of a batch job.  | 9<br>9           |
| A <i>stage in</i> batch request is generated and sent to the server over the connection specified by <i>connect</i> which is the return value of <b>pbs_connect()</b> .  | 9<br>9           |
| 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.  | 9<br>9           |
| The argument, <i>job_id</i> , identifies which job for which file staging is to begin. It is specified in the form: <i>sequence_number.server</i>  | 9<br>9           |
| The argument, <i>location</i> , 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 server does not understand the location, it will reject the request and return an error. If the job is then directed to run at different location, the run request will be rejected. | 9<br>9<br>9<br>9 |
| The argument, <i>extend</i> , is reserved for implementation defined extensions. It is not currently used by this function.  | 9<br>9           |
| <b>SEE ALSO</b>  | 9                |
| qrun(8B), qsub(1B), and pbs_connect(3B)  | 9                |
| <b>DIAGNOSTICS</b>   | 9                |
| When the batch request generated by <b>pbs_stagein()</b> 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 <i>pbs_errno</i> .   | 9<br>9<br>9      |

## 4.2.19. Status Job 9

**NAME** 9

pbs\_statjob - obtain status of pbs batch jobs 9

**SYNOPSIS** 9

#include &lt;pbs\_error.h&gt; 9

#include &lt;pbs\_ifl.h&gt; 9

**struct batch\_status \*pbs\_statjob(int connect, char \*id,  
struct attrl \*attrib, char \*extend)** 9

**void pbs\_statfree(struct batch\_status \*psj)** 9

**DESCRIPTION** 9

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 the string starts with a number, it is a job identifier. If it starts with a alphabetic character, it is 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 a attribute name as listed in *pbs\_alter(3)* and *pbs\_submit(3)*. The *resource* member is only used if the name member is *ATTR\_L*, 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. It is not currently used by this function.

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;
};

```



|  |   |
|--|---|
| }  | 9 |
| It is up to the user to free the structure when no longer needed, by calling <b>pbs_statfree()</b> .   | 9 |
| <b>SEE ALSO</b>  | 9 |
| qstat(1B) and pbs_connect(3B)  | 9 |
| <b>DIAGNOSTICS</b>   | 9 |
| When the batch request generated by <b>pbs_statjob()</b> 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 <code>batch_status</code> structures. If no jobs were available to query or an error occurred, a null pointer is returned. The global integer <code>pbs_errno</code> should be examined to determine the cause. | 9 |

## 4.2.20. Status Queue 9

**NAME** 9

pbs\_statque - obtain status of pbs batch queues 9

**SYNOPSIS** 9

#include &lt;pbs\_error.h&gt; 9

#include &lt;pbs\_ifl.h&gt; 9

**struct batch\_status \*pbs\_statque(int connect, char \*id,** 9**struct attrl \*attrib,** 9**char \*extend)** 9**void pbs\_statfree(struct batch\_status \*psj)** 9**DESCRIPTION** 9

Issue a batch request to obtain the status of a batch queue. 9

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()**. 9The *id* is the name of a queue, in the form: 9

queue\_name 9

or the null string. If *queue\_name* is specified, the status of the queue named *queue\_name* at the server will be returned. If the *id* is a null string or null pointer, the status of all queues at the server will be returned. 9The parameter, *attrib*, is a pointer to an *attrl* structure which is defined in *pbs\_ifl.h* as: 9

```

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

```

9
9
9
9
9
9

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 a 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. 9

The parameter, *extend*, is reserved for implementation defined extensions. It is not currently used by this function. 9

The return value is a pointer to a list of *batch\_status* structures, which is defined in *pbs\_ifl.h* as: 9

```

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

```

9
9
9
9
9
9

It is up to the user to free the structure when no longer needed, by calling **pbs\_statfree()**. 9

|   |   |
|---|---|
| <b>SEE ALSO</b>   | 9 |
| qstat(1B) and pbs_connect(3B)   | 9 |
| <b>DIAGNOSTICS</b>  | 9 |
| When the batch request generated by <b>pbs_statque()</b> function has been completed successfully by a batch server, the routine will return a pointer to the <code>batch_status</code> structure. Otherwise, a null pointer is returned and the error code is set in the global integer <code>pbs_errno</code> . | 9 |

## 4.2.21. Status Server 9

**NAME** 9

pbs\_statserver - obtain status of a pbs batch server 9

**SYNOPSIS** 9

#include &lt;pbs\_error.h&gt; 9

#include &lt;pbs\_ifl.h&gt; 9

**struct batch\_status \*pbs\_statserver(int/ connect, struct/ attrl/ \*attrib, char/ \*extend)** 9

**void pbs\_statfree(struct batch\_status \*psj)** 9

**DESCRIPTION** 9

Issue a batch request to obtain the status of a batch server. 9

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()**. 9

The parameter, *attrib*, is a pointer to an *attrl* structure which is defined in pbs\_ifl.h as: 9

```

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\_alter(3) and pbs\_submit(3). The *resource* member is only used if the *name* member is ATTR\_L, otherwise it should be a pointer to a null string. The *value* member should always be a pointer to a null string. 9

The parameter, *extend*, is reserved for implementation defined extensions. It is not currently used by this function. 9

The return value is a pointer to a list of batch\_status structures, which is defined in pbs\_ifl.h as: 9

```

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

```

It is up to the user to free the space when no longer needed, by calling **pbs\_statfree()**. 9

**SEE ALSO** 9

qstat(1B) and pbs\_connect(3B) 9

**DIAGNOSTICS** 9

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. 9

Otherwise, a null pointer is returned and the error code is set in `pbs_errno`.

9

## 4.2.22. Submit Job 9

**NAME** 9

pbs\_submit - submit a pbs batch job 9

**SYNOPSIS** 9

#include &lt;pbs\_error.h&gt; 9

#include &lt;pbs\_ifl.h&gt; 9

**char \*pbs\_submit(int connect, struct attrpl \*attrib,** 9**char \*script, char \*destination, char \*extend)** 9**DESCRIPTION** 9

Issue a batch request to submit a new batch job. 9

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*. 9

The parameter, *attrib*, is a list of *attrpl* structures which is defined in pbs\_ifl.h as: 9

```

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

```

9
9
9
9
9
9
9

The *attrib* list is terminated by the first entry where *next* is a null pointer. 9

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: 9

```

#define ATTR_a "Execution_Time"
    Defines the job's execution time.
#define ATTR_A "Account_Name"
    Defines the account string.
#define ATTR_c "Checkpoint"
    Defines the checkpoint interval.
#define ATTR_e "Error_Path"
    Defines the path name for the standard error of the job.
#define ATTR_g "Group_List"
    Defines the list of group names under which the job may execute.
#define ATTR_h "Hold_Types"
    Defines the hold types, the only allowable value string is "u".
#define ATTR_j "Join_Paths"
    Defines whether standard error and standard output are joined (merged).
#define ATTR_k "Keep_Files"
    Defines which output of the job is kept on the execution host.
#define ATTR_l "Resource_List"
    Defines a resource required by the job.

```

9
9
9
9
9
9
9
9
9
9
9
9
9
9
9
9
9
9
9
9

```

#define ATTR_m "Mail_Points"          9
    Defines the points at which the server will send mail about the job.  9
#define ATTR_M "Mail_Users"          9
    Defines the list of users who would receive mail about the job.      9
#define ATTR_N "Job_Name"            9
    Defines the job name.                                                9
#define ATTR_o "Output_Path"         9
    Defines the path name for the standard output of the job.           9
#define ATTR_p "Priority"             9
    Defines the priority of the job.                                     9
#define ATTR_r "Rerunable"           9
    Defines the rerunable flag.                                          9
#define ATTR_S "Shell_Path_List"     9
    Defines the path to the shell which will interpret the job script.   9
#define ATTR_u "User_List"           9
    Defines the list of user names under which the job may execute.      9
#define ATTR_v "Variable_List"       9
    Defines the list of additional environment variables which are exported to the job. 9
#define ATTR_depend "depend"         9
    Defines the inter-job dependencies.                                  9
#define ATTR_stagein "stagein"       9
    Defines the list of files to be staged in prior to job execution.    9
#define ATTR_stageout "stageout"     9
    Defines the list of files to be staged out after job execution.      9

```

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. 9

Associated with an attribute of type ATTR\_l (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*. 9

The *op* member is forced to a value of SET by `pbs_submit()`. 9

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. 9

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. 9

The parameter, *extend*, is reserved for implementation defined extensions. It is not currently used by this function. 9

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. 9

**SEE ALSO** 9  
`qsub(1B)` and `pbs_connect(3B)` 9

**DIAGNOSTICS** 9  
When the batch request generated by `pbs_submit()` function has been completed successfully 9

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`.

9  
9  
9



|  |   |
|--|---|
| 4.2.23. Terminate Server   | 9 |
| <b>NAME</b>  | 9 |
| pbs_terminate - terminate a pbs batch server   | 9 |
| <b>SYNOPSIS</b>  | 9 |
| #include <pbs_error.h>   | 9 |
| #include <pbs_ifl.h>   | 9 |
| <b>int pbs_terminate(int connect, int manner, char *extend)</b>  | 9 |
| <b>DESCRIPTION</b>   | 9 |
| Issue a batch request to shut down a batch server. This request requires the privilege level usually reserved for batch operators and administrators.  | 9 |
| A <i>Server Shutdown</i> batch request is generated and sent to the server over the connection specified by <i>connect</i> which is the return value of <b>pbs_connect()</b> .   | 9 |
| The parameter, <i>manner</i> , specifies the manner in which the server is shut down. The available manners are defined in pbs_ifl.h as:   | 9 |
| #define SHUT_IMMEDIATE 0   | 9 |
| Shutdown is to be immediate, runnings jobs are checkpointed, requeued, or deleted as required.   | 9 |
| #define SHUT_DELAY 1   | 9 |
| Jobs which can be checkpointed are checkpointed, terminated, and requeued.   | 9 |
| Jobs which cannot be checkpointed but are rerunnable are terminated and requeued. Shutdown is delayed until the remaining running jobs complete. No new jobs will be started by the server.  | 9 |
| #define SHUT_QUICK 2   | 9 |
| Shutdown of the server occurs as soon as the server can record lastest state. Jobs which are currently running, are left in the Running state.   | 9 |
| The server will not respond to the batch request until the server has completed its termination procedure.   | 9 |
| The parameter, <i>extend</i> , is reserved for implementation defined extensions. It is not currently used by this function.   | 9 |
| This call requires PBS Operator or Manager privilege.  | 9 |
| <b>SEE ALSO</b>  | 9 |
| qterm(8B) and pbs_connect(3B)  | 9 |
| <b>DIAGNOSTICS</b>   | 9 |
| When the batch request generated by <b>pbs_terminate()</b> 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. | 9 |

## 4.2.24. Query Resource Availability 9

While the following functions may be used by any user, they are intended for use in a batch scheduler. They replace `avail`, `totpool`, and `usepool` calls made to `pbs_mom` in early versions of PBS. 9

**NAME** 9

`pbs_resquery`, `avail`, `totpool`, `usepool` - query resource availability 9

**SYNOPSIS** 9

```
#include <pbs_error.h> 9
```

```
#include <pbs_ifl.h> 9
```

```
int pbs_resquery(int connect, char **resourcelist, int *arraysize, 9
int *available, int *allocated, int *reserved, int *down ) 9
```

```
char *avail(int connect, char *resc) 9
```

```
int totpool(int connect, int update) 9
```

```
int usepool(int connect, int update) 9
```

**DESCRIPTION** 9**pbs\_resquery** 9

Issue a request to the batch server to query the availability of 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*. *available*, *allocated*, *reserved*, and *down* are integer arrays of size *arraysize*. The amount of resource specified in the corresponding *resourcelist* string which is available, already allocated, reserved, and down/off-line is returned in the integer arrays. 9

At the present time the only resources which may be specified is "nodes". It may be specified as 9

```
nodes 9
```

```
nodes= 9
```

```
nodes=specification 9
```

where *specification* is what a user specifies in the `-l` option argument list for `nodes`, see `qsub(1B)` and the various `pbs_resource_*` man pages. 9

Where the node *resourcelist* is a simple type, such as "nodes", "nodes=", or "nodes=*type*", the numbers returned reflect the actual number of nodes (of the specified type) which are *available*, *allocated*, *reserved*, or *down*. 9

For a more complex node *resourcelist*, such as "nodes=2" or "nodes=type1:type2", only the value returned in *available* has meaning. If the number in *available* is positive, it is the number of nodes required to satisfied the specification and that some set of nodes are available which will satisfy it, see `avail()`. If the number in *available* is zero, some number of nodes required to satisfied the specification are currently unavailable, the request might be satisfied at a later time. If the number in *available* is negative, no combination of known nodes can satisfied the specification. 9

**avail** 9

The `avail()` call is provided as conversion aid for scheduler written for early versions of PBS. The `avail()` routine uses `pbs_resquery()` and returns a character string answer. *connect* is the connection returned by `pbs_connect()`. *resc* is a sin- 9

gle *node=specification* specification as discussed above. If the nodes to satisfy the specification are currently available, the return value is the character string **yes**. If the nodes are currently unavailable, the return is the character string **no**. If the specification could never be satisfied, the return is the string **never**. An error in the specification returns the character string **?**.

### **totpool** 9

The *totpool()* function returns the total number of nodes known to the PBS server. This is the sum of the number of nodes available, allocated, reserved, and down. The parameter *connection* is the connection returned by *pbs\_connect()*. The parameter *update* if non-zero, causes *totpool()* to issue a *pbs\_resquery()* call to obtain fresh information. If zero, numbers from the prior *pbs\_resquery()* are used. 9

### **usepool** 9

*usepool()* returns the number of nodes currently in use, the sum of allocated, reserved, and down. The parameter *connection* is the connection returned by *pbs\_connect()*. The parameter *update* if non-zero, causes *totpool()* to issue a *pbs\_resquery()* call to obtain fresh information. If zero, numbers from the prior *pbs\_resquery()* are used. 9

### **SEE ALSO** 9

*qsub(1B)*, *pbs\_connect(3B)*, *pbs\_disconnect(3B)*, *pbs\_resreserve(3B)* and *pbs\_resources(7B)* 9

### **DIAGNOSTICS** 9

When the batch request generated by the ***pbs\_resquery()*** 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*. 9

The functions *usepool()* and *totpool()* return -1 on error. 9

|  |   |
|--|---|
| 4.2.25. Make and Release Resource Reservations   | 9 |
| The following functions require manager or operator privilege. They are intended for use in a batch scheduler.   | 9 |
| <b>NAME</b>  | 9 |
| <code>pbs_resreserve, pbs_resrelease</code> - reserve/free batch resources   | 9 |
| <b>SYNOPSIS</b>  | 9 |
| <code>#include &lt;pbs_error.h&gt;</code>  | 9 |
| <code>#include &lt;pbs_ifl.h&gt;</code>  | 9 |
| <b>int pbs_resreserve(int connect, char **resourcelist, int arraysize, resource_t *resource_id)</b>  | 9 |
| <b>int pbs_resrelease(int connect, resource_t resource_id)</b>   | 9 |
| <b>DESCRIPTION</b>   | 9 |
| <b>pbs_resreserver</b>   | 9 |
| Issue a request to the batch server to reserve specified resources. <i>connect</i> is the connection returned by <b>pbs_connect()</b> . <i>resourcelist</i> is an array of one or more strings specifying the resources to be queried. <i>arraysize</i> is the is the number of strings in <i>resourcelist</i> . <i>resource_id</i> is a pointer to a resource handle. The pointer cannot be null. If the present value of the resource handle is <b>RESOURCE_T_NULL</b> , this request is for a new reservation and if successful, a resource handle will be returned in <i>resource_id</i> . | 9 |
| If the value of <i>resource_id</i> as supplied by the caller is not <b>RESOURCE_T_NULL</b> , this is a 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 <i>pbs_resrelease()</i> .   | 9 |
| At the present time the only resources which may be specified are "nodes". It should be specified as <code>nodes=specification</code> where <i>specification</i> is what a user specifies in the <code>-l</code> option argument list for nodes, see <i>qsub</i> (1B).   | 9 |
| <b>pbs_resrelease</b>  | 9 |
| The <i>pbs_resrelease()</i> call releases or frees resources reserved with the resource handle of <i>resource_id</i> returned from a prior <i>pbs_resreserve()</i> call. <i>connect</i> is the connection returned by <b>pbs_connect()</b> .   | 9 |
| Both functions require that the issuing user have operator or administrator privilege.   | 9 |
| <b>SEE ALSO</b>  | 9 |
| <i>qsub</i> (1B), <i>pbs_connect</i> (3B), <i>pbs_disconnect</i> (3B), <i>pbs_resquery</i> (3B) and <i>pbs_resources</i> (7B)  | 9 |
| <b>DIAGNOSTICS</b>   | 9 |
| <i>pbs_resreserve()</i> and <i>pbs_resrelease()</i> return zero on success. Otherwise, a non zero error is returned. The error number is also set in <i>pbs_errno</i> .  | 9 |
| <b>PBSE_RMPART</b>   | 9 |
| 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.  | 9 |
| <b>PBSE_RMBADPARAM</b>   | 9 |
| a parameter is incorrect, such as a null for the pointer to the <i>resource_id</i> .   | 9 |

PBSE\_RMNOPARAM 9  
 a parameter is missing, such as a null resource list. 9

#### 4.2.26. Node Status 9

The following function provides a means to query the status of PBS nodes. A node has a name and two attributes. ATTR\_NODE\_state (state) may take the values of *free* or *one* or more of the following "or-ed" together: 9

*free* The node is not being used by PBS and is free to be allocated to a job. 9

*down* The node is not responding to queries by the PBS job server. 9

*offline*  
 The node has been marked off-line by the administrator 9

*reserve*  
 The node has been reserved by the PBS job scheduler. 9

*job-exclusive*  
 The node has been assigned exclusively to a job, no other jobs can run on that node. 9

*job-sharing*  
 The node has been assigned to one or more jobs. Other jobs willing to share nodes may be allocated to it. 9

The attribute ATTR\_NODE\_properties (properties) is a list of alternative names which may be used to identify sets or groups of nodes by some arbitrary criteria. 9

#### NAME 9

`pbs_statnode` - obtain status of pbs nodes 9

#### SYNOPSIS 9

```
#include <pbs_error.h> 9
```

```
#include <pbs_ifl.h> 9
```

```
struct batch_status *pbs_statnode(int connect, char *id, struct attrl *attrib,  
char *extend) 9
```

```
void pbs_statfree(struct batch_status *psj) 9
```

#### DESCRIPTION 9

Issue a batch request to obtain the status of a PBS node or nodes. 9

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()`. 9

The *id* is the name of a node or the null string. If *id* specifies a node name, the status of that node will be returned. If the *id* is a null string (or null pointer), the status of all nodes at the server will be returned. 9

The parameter, *attrib*, is a pointer to an *attrl* structure which is defined in `pbs_ifl.h` as: 9

```
struct attrl { 9  
    struct attrl *next; 9  
    char *name; 9  
    char *resource; 9  
    char *value; 9  
}; 9
```

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 only supported attribute names relating to nodes are "state" and "properties". 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. It is not currently used by this function.

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()**.

#### SEE ALSO

qstat(1B) and pbs\_connect(3B)

#### DIAGNOSTICS

When the batch request generated by **pbs\_statnode()** 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\_errno*.

[Page intentionally left blank.]

|  |                       |
|--|-----------------------|
|  | 9                     |
| <b>5. User Commands</b>  | <b>9</b>              |
| This section describes the commands available to the general user. Unless otherwise noted, the command must conform to the POSIX 1003.2d specification of the command as to syntax and functionality.  | 9<br>9<br>9           |
| <b>5.1. General Specifications of User Commands</b>  | <b>9</b>              |
| The following specifications apply to all user commands.   | 9                     |
| <b>5.1.1. Error Checking</b>   | <b>9</b>              |
| All user commands will validate their invocation for the correct syntax. Any syntax error or invalid option will terminate the command with an error message on standard error and an exit status greater than zero.   | 9<br>9<br>9           |
| <b>5.1.2. Directing Requests to Correct Server</b>   | <b>9</b>              |
| A command performs its function by sending the corresponding request for service to a batch server. The choice of batch servers to which to send the request is governed by the following ordered set of rules:  | 9<br>9<br>9           |
| 1. For those commands which require or accept a job identifier operand, if the server is specified in the job identifier operand as <code>@server</code> , then the batch requests will be sent to the server named by <code>server</code> .   | 9<br>9<br>9           |
| 2. For those commands which require or accept a job identifier operand and the <code>@server</code> is not specified, then the command will attempt to determine the current location of the job by sending a <i>Locate Job</i> batch request to the server which created the job.   | 9<br>9<br>9           |
| 3. If a server component of a destination is supplied via the <code>-q</code> option, such as on <b>qsub</b> and <b>qselect</b> , but not <b>qalter</b> , then the server request is sent to that server.  | 9<br>9                |
| 4. The server request is sent to the server identified as the default server, see section 2.7.4.   | 9                     |
| <b>5.1.3. Operands</b>   | <b>9</b>              |
| Unless noted, when more than one operand is specified on the command line, the command processes each operand in turn. An error reply from a server on one operand will be noted in the standard error stream. The command continues processing the other operands. If an error reply was received for any operand, the final exit status for the command will be greater than zero. | 9<br>9<br>9<br>9<br>9 |
| Generally, the operands to commands will be job identifiers as described in section 2.7.6 or destination identifiers described in section 2.7.3.   | 9<br>9                |
| <b>5.2. General User Commands</b>  | <b>9</b>              |



## 5.2.1. Alter Job 9

**NAME** 9

qalter – alter pbs batch job 9

**SYNOPSIS** 9

qalter [-a *date\_time*] [-A *account\_string*] [-c *interval*] [-e *path*] [-h *hold\_list*] [-j *join*] [-k *keep*] [-l *resource\_list*] [-m *mail\_options*] [-M *user\_list*] [-N *name*] [-o *path*] [-p *priority*] [-r *c*] [-S *path*] [-u *user\_list*] [-W *additional\_attributes*] *job\_identifier*...

**DESCRIPTION** 9

The **qalter** command modifies the attributes of the job or jobs specified by *job\_identifier* on the command line. Only those attributes listed as options on the command will be modified. If any of the specified attributes cannot be modified for a job for any reason, none of that job's attributes will be modified.

The qalter command accomplishes the modifications by sending a *Modify Job* batch request to the batch server which owns each job.

**OPTIONS** 9-a *date\_time* 9

Replaces the attribute Execution\_Time time at which the job becomes eligible for execution. The *date\_time* argument syntax is: *[[[CC]YY]MM]DD]hhmm[.SS]*.

If the month, *MM*, is not specified, it will default to the current month if the specified day *DD*, is in the future. Otherwise, the month will be set to next month. Likewise, if the day, *DD*, is not specified, it will default to today if the time *hhmm* is in the future. Otherwise, the day will be set to tomorrow. This *date\_time* will be converted to the integer number of seconds since Epoch that is equivalent to the local time on the system where the command is being executed.

This attribute can be altered once the job has begun execution, but it will not take affect until the job is rerun.

-A *account\_string* 9

Replaces the the Account\_Name attribute, the account string associated with the job. The syntax of the *account\_string* is defined in section 2.7.1. It is interpreted by the server which executes the job.

This attribute cannot be altered once the job has begun execution.

-c *interval* 9

Replaces the Checkpoint attribute, the interval at which the job will be checkpointed. If the job executes upon a host which does not support checkpoint, this option will be ignored.

The *interval* argument is specified as:

n No checkpointing is to be performed. The job's Checkpoint attribute is set to the string "n".

s Checkpointing is to be performed only when the server executing the job is shut-down. The job's Checkpoint attribute is set to the string "s".

c Checkpointing is to be performed at the default minimum cpu time for the queue from which the job is executing. The job's Checkpoint attribute is set to the string "c".

c=*minutes* 9

Checkpointing is to be performed at an interval of *minutes*, which is the integer number of minutes of CPU time used by the job. This value must be greater than zero. If the number is less than the default checkpoint time, the default time will

- be used. The Checkpoint attributes is set to the string specified by "*c=minutes*". 9
- This attribute can be altered once the job has begun execution, but the new value 9  
does not take affect until the job is rerun. 9
- e path** Replaces the Error\_Path attribute, the path to be used for the standard error stream 9  
of the batch job. The *path* argument is of the form: 9  
*[hostname:]path\_name* 9  
where *hostname* is the name of a host to which the file will be returned and 9  
*path\_name* is the path name on that host in the syntax recognized by POSIX 9  
1003.1. The argument will be interpreted as follows: 9
- path\_name* 9  
Where *path\_name* is not an absolute path name, then the qalter command will 9  
expand the path name relative to the current working directory of the com- 9  
mand. The command will supply the name of the host upon which it is execut- 9  
ing for the *hostname* component. 9
- hostname:path\_name* 9  
Where *path\_name* is not an absolute path name, then the qalter command will 9  
not expand the path name. The execution server will expand it relative to the 9  
home directory of the user on the system specified by *hostname*. 9
- path\_name* 9  
Where *path\_name* specifies an absolute path name, then qalter will supply the 9  
name of the host on which it is executing for the *hostname*. 9
- hostname:path\_name* 9  
Where *path\_name* specifies an absolute path name, the path will be used as 9  
specified. 9
- This attribute can be altered once the job has begun execution, but it will not take 9  
affect until the job is rerun. 9
- h hold\_list** 9  
Updates the Hold\_Types attribute, the types of holds on the job. The *hold\_list* argu- 9  
ment is a string of one or more of the following characters: 9
- u** Add the USER type hold. 9
- s** Add the SYSTEM type hold if the user has the appropriate level of privilege. 9  
[Typically reserved to the batch administrator.] 9
- o** Add the OTHER (or OPERATOR ) type hold if the user has the appropriate 9  
level of privilege. [Typically reserved to the batch administrator and batch 9  
operator.] 9
- n** Set to none; that is clear the hold types which could be applied with the users 9  
level of privilege. 9
- Repetition of characters is permitted, but "n" may not appear in the same option ar- 9  
gument with the other three characters. This attribute can be altered once the job 9  
has begun execution, but the hold will not take affect until the job is rerun. 9
- j join** Declares which standard streams of the job will be merged together. The *join* argu- 9  
ment value may be the characters "oe" and "eo", or the single character "n". 9
- A argument value of *oe* directs that the standard output and standard error streams 9  
of the job will be merged, intermixed, and returned as the standard output. A argu- 9  
ment value of *eo* directs that the standard output and standard error streams of the 9  
job will be merged, intermixed, and returned as the standard error. The Join\_Path 9  
job attribute is set to the value. 9
- A value of *n* directs that the two streams will be two separate files. The Join\_Path at- 9  
tribute is set to "n". This attribute can be altered once the job has begun execution, 9  
but it will not take affect until the job is rerun. 9

- k keep** Defines which if either of standard output or standard error of the job will be retained on the execution host. If set for a stream, this option overrides the path name for that stream. 9
- The argument is either the single letter "e", "o", or "n", or one or more of the letters "e" and "o" combined in either order. 9
- n** No streams are to be retained. The `Keep_Files` attribute is set to `KEEP_NONE`, "n". 9
  - e** The standard error stream is to be retained on the execution host. The stream will be placed in the home directory of the user under whose user id the job executed. The file name will be the default file name given by: `job_name.e``sequence` where `job_name` is the name specified for the job, and `sequence` is the sequence number component of the job identifier. The attribute is set to include `KEEP_ERROR`, "e". 9
  - o** The standard output stream is to be retained on the execution host. The stream will be placed in the home directory of the user under whose user id the job executed. The file name will be the default file name given by: `job_name.o``sequence` where `job_name` is the name specified for the job, and `sequence` is the sequence number component of the job identifier. The `Output_Path` attribute is set to include `KEEP_OUTPUT`, "o". 9
  - eo** Both the standard output and standard error streams will be retained. The attribute is set to `KEEP_OUTPUT | KEEP_ERROR`. 9
  - oe** Both the standard output and standard error streams will be retained. The attribute is set to `KEEP_OUTPUT | KEEP_ERROR`. 9
- Repetition of characters is permitted, but "n" may not appear in the same option argument with the other two characters. This attribute cannot be altered once the job has begun execution. 9
- l resource\_list** 9
- Modifies the `Resource_List` attribute, the list of resources that are required by the job. The `Resource_List` argument is in the following syntax: 9
- ```
resource_name[=value][,resource_name[=value]],...
```
- For each resource listed, if a resource with the specified name already exist in the jobs resource attribute, the value for that resource will be updated. If the named resource does not exist in the job resource attribute, the resource name and value will be added. No white space is allowed in the value. 9
- Because the list of supported resources vary from host to host, the command will perform no validation of the name or value. 9
- If a requested modification to a resource would exceed the resource limits for jobs in the current queue, the server will reject the request. 9
- If the job is running, only certain, resources can be altered. Which resources can be altered in the run state is system dependent. A user may only lower the limit for those resources. A PBS Manager or Operator may increase them. 9
- m mail\_options** 9
- Replaces the set of conditions under which the execution server will send a mail message about the job. The `mail_options` argument is a string which consists of one or more repetitions of the single character "n", or one or more repetitions of the characters "a", "b", and "e". 9
- If the character "n" is specified, no mail will be sent. The `Mail_Points` attribute is set to `NONE`, "n". 9
- For the letters "a", "b", and "e": 9

- a mail is sent when the job is aborted by the batch system. The Mail\_Points attribute is set to ABORT, "a". 9  
9
- b mail is sent when the job begins execution. The Mail\_Points attribute is set to BEGINNING, "b". 9  
9
- e mail is sent when the job terminates. The Mail\_Points attribute is set to EXIT, "e". 9  
9
- M *user\_list* 9  
Replaces the list of users to whom mail is sent by the execution server when it sends mail about the job. 9  
9  
The *user\_list* argument is of the form: 9  
*user[@host][,user[@host],...]* 9  
The Mail\_Users attribute is set to the argument. 9
- N *name* Renames the job. The name specified may be up to and including 15 characters in length. It must consist of printable, non white space characters with the first character alphabetic. [See the discussion of the -N option under *qsub(1)*.] The Job\_Name attribute is reset to the name value. 9  
9
- o *path* Replaces the path to be used for the standard output stream of the batch job. The *path* argument is of the form: 9  
9  
*[hostname:]path\_name* 9  
where *hostname* is the name of a host to which the file will be returned and *path\_name* is the path name on that host in the syntax recognized by POSIX. The argument will be interpreted as follows: 9  
9
- path\_name* 9  
Where *path\_name* is not an absolute path name, then the *qalter* command will expand the path name relative to the current working directory of the command. The command will supply the name of the host upon which it is executing for the *hostname* component. 9  
9
- hostname:path\_name* 9  
Where *path\_name* is not an absolute path name, then the *qalter* command will not expand the path name. The execution server will expand it relative to the home directory of the user on the system specified by *hostname*. 9  
9
- path\_name* 9  
Where *path\_name* specifies an absolute path name, then the *qalter* will supply the name of the host on which it is executing for the *hostname*. 9  
9
- hostname:path\_name* 9  
Where *path\_name* specifies an absolute path name, the path will be used as specified. 9  
9
- This attribute can be altered once the job has begun execution, but it will not take affect until the job is rerun. 9  
9
- p *priority* 9  
Replaces the priority of the job. The *priority* argument must be a integer between -1024 and +1023 inclusive. The Priority attribute is set to this signed integer value. 9  
9  
This attribute can be altered once the job has begun execution, but it will not take affect until the job is rerun. 9  
9
- r *c* Declares whether the job is rerunable. See the **qrerun** command. The option argument *c* is a single character. PBS recognizes the following characters: *y* and *n*. Also see *rerunability* in the glossary. 9  
9  
If the argument is "y", the job is marked rerunable. The Rerunable attribute is set to 'y'. If the argument is "n", the job is marked as not rerunable. The Rerunable attribute is set to 'n'. 9  
9

- S path** Declares the shell that interprets the job script. 9
- The option argument *path\_list* is in the form: 9
- path[@host][,path[@host],...]* 9
- Only one path may be specified for any host named. Only one path may be specified 9
- without the corresponding host name. The path selected will be the one with the 9
- host name that matched the name of the execution host. If no matching host is 9
- found, then if present the path specified without a host will be selected. 9
- If the **-S** option is not specified, the option argument is the null string, or no entry 9
- from the *path\_list* is selected, the execution will use the login shell of the user on 9
- the execution host. The Shell\_Path\_List attribute is set to the *path\_list* argument if 9
- present, otherwise it is set to the null string. 9
- This attribute can be altered once the job has begun execution, but it will not take 9
- affect until the job is rerun. 9
- u user\_list** 9
- Replaces the user name under which the job is to run on the execution system. 9
- The *user\_list* argument is of the form: 9
- user[@host][,user[@host],...]* 9
- Only one user name may be given for per specified host. Only one of the *user* speci- 9
- fications may be supplied without the corresponding *host* specification. That user 9
- name will be used for execution on any host not named in the argument list. The 9
- User\_List attribute is set to the value of *User\_List*. 9
- This attribute cannot be altered once the job has begun execution. 9
- W additional\_attributes** 9
- The **-W** option allows for the modification of additional job attributes. The general 9
- syntax of the **-W** is in the form: 9
- W attr\_name=value[,attr\_name=value...]** 9
- Note if white space occurs anywhere within the option argument string or the equal 9
- sign, "=", occurs within an *attribute\_value* string, then the string must be enclosed 9
- with either single or double quote marks. 9
- PBS currently supports the following attributes within the **-W** option. 9
- depend=dependency\_list* 9
- Redefines the depend attribute listing the dependencies between this and other jobs. 9
- The *dependency\_list* is in the form: *type[:argument[:argument...]][,type:argument...]*. 9
- The *argument* is either a numeric count or a PBS job id according to *type*. If argu- 9
- ment is a count, it must be greater than 0. If it is a job id and is not fully specified 9
- in the form: *seq\_number.server.name*, it will be expanded according to the default 9
- server rules. If *argument* is null (the preceding colon need not be specified), the de- 9
- pendency of the cooresponding type is cleared (unset). 9
- synccount:count* 9
- This job is the first in a set of jobs to be executed at the same time. *Count* is the 9
- number of additional jobs in the set. 9
- syncwith:jobid* 9
- This job is an additional member of a set of jobs to be executed at the same time. 9
- Jobid* is the job identifier of the first job in the set. 9
- after:jobid[:jobid...]* 9
- This job may be scheduled for execution at any point after jobs *jobid* have start- 9
- ed execution. 9
- afterok:jobid[:jobid...]* 9
- This job may be scheduled for execution only after jobs *jobid* have terminated 9
- with no errors. 9

|                                                                                                                                                                                                                                                                                                                                                         |   |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| <i>afternotok: jobid[:jobid...]</i>                                                                                                                                                                                                                                                                                                                     | 9 |
| This job may be scheduled for execution only after jobs <i>jobid</i> have terminated with errors.                                                                                                                                                                                                                                                       | 9 |
| <i>afterany: jobid[:jobid...]</i>                                                                                                                                                                                                                                                                                                                       | 9 |
| This job may be scheduled for execution after jobs <i>jobid</i> have terminated, with or without errors.                                                                                                                                                                                                                                                | 9 |
| <i>on: count</i>                                                                                                                                                                                                                                                                                                                                        | 9 |
| This job may be scheduled for execution after <i>count</i> dependencies on other jobs have been satisfied. This form is used in conjunction with one of the <i>before</i> forms, see below.                                                                                                                                                             | 9 |
| <i>before: jobid[:jobid...]</i>                                                                                                                                                                                                                                                                                                                         | 9 |
| When this job has begun execution, then jobs <i>jobid...</i> may begin.                                                                                                                                                                                                                                                                                 | 9 |
| <i>beforeok: jobid[:jobid...]</i>                                                                                                                                                                                                                                                                                                                       | 9 |
| If this job terminates execution without errors, then jobs <i>jobid...</i> may begin.                                                                                                                                                                                                                                                                   | 9 |
| <i>beforenotok: jobid[:jobid...]</i>                                                                                                                                                                                                                                                                                                                    | 9 |
| If this job terminates execution with errors, then jobs <i>jobid...</i> may begin.                                                                                                                                                                                                                                                                      | 9 |
| <i>beforeany: jobid[:jobid...]</i>                                                                                                                                                                                                                                                                                                                      | 9 |
| When this job terminates execution, jobs <i>jobid...</i> may begin.                                                                                                                                                                                                                                                                                     | 9 |
| If any of the <i>before</i> forms are used, the job referenced by <i>jobid</i> must have been submitted with a dependency type of <i>on</i> .                                                                                                                                                                                                           | 9 |
| The job specified in any of the <i>before</i> forms must have the same owner as the job being altered.. Otherwise, the dependency will not take effect.                                                                                                                                                                                                 | 9 |
| Error processing of the existence, state, or condition of the job on which the newly submitted job is a deferred service, i.e. the check is performed after the job is queued. If an error is detected, the new job will be deleted by the server. Mail will be sent to the job submitter stating the error.                                            | 9 |
| <i>group_list=g_list</i>                                                                                                                                                                                                                                                                                                                                | 9 |
| Alters the <i>group_list</i> attribute, which lists the group name under which the job is to run on the execution system.                                                                                                                                                                                                                               | 9 |
| The <i>g_list</i> argument is of the form: <i>group[@host][,group[@host],...]</i>                                                                                                                                                                                                                                                                       | 9 |
| Only one group name may be given per specified host. Only one of the <i>group</i> specifications may be supplied without the corresponding <i>host</i> specification. That group name will be used for execution on any host not named in the argument list.                                                                                            | 9 |
| <i>stagein=file_list</i>                                                                                                                                                                                                                                                                                                                                | 9 |
| <i>stageout=file_list</i>                                                                                                                                                                                                                                                                                                                               | 9 |
| Alters the <i>stageout</i> attribute or the <i>stagein</i> attribute, which list which files are staged (copied) in before job start or staged out after the job completes execution. The <i>file_list</i> is in the form: <i>local_file@hostname:remote_file[,...]</i>                                                                                 | 9 |
| The name <i>local_file</i> is the name on the system where the job executes. It may be an absolute path or a path relative to the home directory of the user. The name <i>remote_file</i> is the destination name on the host specified by <i>hostname</i> . The name may be absolute or relative to the user's home directory on the destination host. | 9 |

**OPERANDS**

The *qalter* command accepts one or more *job\_identifier* operands of the form:  
*sequence\_number[.server\_name][@server]*

**STANDARD ERROR**

Any error condition, either in processing the options or the operands, or any error received in reply to the batch requests will result in a error message being written to standard error.

|                                                                                                                                               |   |
|-----------------------------------------------------------------------------------------------------------------------------------------------|---|
| <b>EXIT STATUS</b>                                                                                                                            | 9 |
| Upon successful processing of all the operands presented to the the qalter command, the exit status will be a value of zero.                  | 9 |
| If the qalter command fails to process any operand, the command exits with a value greater than zero.                                         | 9 |
| <b>SEE ALSO</b>                                                                                                                               | 9 |
| qsub(1B), qstat(1B), pbs_alterjob(3B), pbs_statjob(3B), pbs_selectjob(3B), pbs_re-<br>sources_*(7B), where * is system type, and the PBS ERS. | 9 |

|                                                                                                                                                                                                                                                                                                                                                                           |   |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| 5.2.2. Delete Job                                                                                                                                                                                                                                                                                                                                                         | 9 |
| <b>NAME</b>                                                                                                                                                                                                                                                                                                                                                               | 9 |
| qdel - delete pbs batch job                                                                                                                                                                                                                                                                                                                                               | 9 |
| <b>SYNOPSIS</b>                                                                                                                                                                                                                                                                                                                                                           | 9 |
| qdel [-W delay   force] job_identifier ...                                                                                                                                                                                                                                                                                                                                | 9 |
| <b>DESCRIPTION</b>                                                                                                                                                                                                                                                                                                                                                        | 9 |
| The <b>qdel</b> command deletes jobs in the order in which their job identifiers are presented to the command. A job is deleted by sending a <i>Delete Job</i> batch request to the batch server that owns the job. A job that has been deleted is no longer subject to management by batch services.                                                                     | 9 |
| A batch job may be deleted by its owner, the batch operator, or the batch administrator.                                                                                                                                                                                                                                                                                  | 9 |
| A batch job being deleted by a server will be sent a <b>SIGTERM</b> signal following by a <b>SIGKILL</b> signal. The time delay between the two signals is an attribute of the execution queue from which the job was run (settable by the administrator). This delay may be overridden by the <i>-W delay</i> option, where <i>delay</i> is a time specified in seconds. | 9 |
| The <i>-W force</i> option, where <i>force</i> is the literal character string <i>force</i> , directs that the job is to be deleted even if the node on which the job is executing is unreachable.                                                                                                                                                                        | 9 |
| See the PBS ERS section 3.1.3.3, "Delete Job Request", for more information.                                                                                                                                                                                                                                                                                              | 9 |
| <b>OPTIONS</b>                                                                                                                                                                                                                                                                                                                                                            | 9 |
| <i>-W delay</i> Specify the delay between the sending of the SIGTERM and SIGKILL signals. The argument <i>delay</i> specifies a unsigned integer number of seconds.                                                                                                                                                                                                       | 9 |
| <b>OPERANDS</b>                                                                                                                                                                                                                                                                                                                                                           | 9 |
| The qdel command accepts one or more <i>job_identifier</i> operands of the form:<br><i>sequence_number[.server_name][@server]</i>                                                                                                                                                                                                                                         | 9 |
| See the description under "Job Identifier" in section 2.7.6 in this ERS.                                                                                                                                                                                                                                                                                                  | 9 |
| <b>STANDARD ERROR</b>                                                                                                                                                                                                                                                                                                                                                     | 9 |
| The qdel command will write a diagnostic messages to standard error for each error occurrence.                                                                                                                                                                                                                                                                            | 9 |
| <b>EXIT STATUS</b>                                                                                                                                                                                                                                                                                                                                                        | 9 |
| Upon successful processing of all the operands presented to the qdel command, the exit status will be a value of zero.                                                                                                                                                                                                                                                    | 9 |
| If the qdel command fails to process any operand, the command exits with a value greater than zero.                                                                                                                                                                                                                                                                       | 9 |
| <b>SEE ALSO</b>                                                                                                                                                                                                                                                                                                                                                           | 9 |
| qsub(1B), qsig(1B), and pbs_deljob(3B)                                                                                                                                                                                                                                                                                                                                    | 9 |



|                                                                                                                                                                                                                                                                                                                                                                                                       |   |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| 5.2.3. Hold Job                                                                                                                                                                                                                                                                                                                                                                                       | 9 |
| <b>NAME</b>                                                                                                                                                                                                                                                                                                                                                                                           | 9 |
| qhold - hold pbs batch jobs                                                                                                                                                                                                                                                                                                                                                                           | 9 |
| <b>SYNOPSIS</b>                                                                                                                                                                                                                                                                                                                                                                                       | 9 |
| qhold [-h hold_list] job_identifier ...                                                                                                                                                                                                                                                                                                                                                               | 9 |
| <b>DESCRIPTION</b>                                                                                                                                                                                                                                                                                                                                                                                    | 9 |
| The <b>qhold</b> command requests that a server place one or more holds on a job. A job that has a hold is not eligible for execution. There are three supported holds: <i>USER</i> , <i>OTHER</i> (also known as operator), and <i>SYSTEM</i> .                                                                                                                                                      | 9 |
| A user may place a <i>USER</i> hold upon any job the user owns. An "operator", who is a user with "operator privilege," may place either an <i>USER</i> or an <i>OTHER</i> hold on any job. The batch administrator may place any hold on any job.                                                                                                                                                    | 9 |
| If no <i>-h</i> option is given, the <i>USER</i> hold will be applied to the jobs described by the <i>job_identifier</i> operand list.                                                                                                                                                                                                                                                                | 9 |
| If the job identified by <i>job_identifier</i> is in the <b>queued</b> , <b>held</b> , or <b>waiting</b> states, then all that occurs is that the hold type is added to the job. The job is then placed into <b>held</b> state if it resides in an execution queue.                                                                                                                                   | 9 |
| If the job is in <b>running</b> state, then the following additional action is taken to interrupt the execution of the job. If checkpoint / restart is supported by the host system, requesting a hold on a running job will (1) cause the job to be checkpointed, (2) the resources assigned to the job will be released, and (3) the job is placed in the <b>held</b> state in the execution queue. | 9 |
| If checkpoint / restart is not supported, qhold will only set the requested hold attribute. This will have no effect unless the job is rerun with the <b>qrerun</b> command.                                                                                                                                                                                                                          | 9 |
| The qhold command sends a <i>Hold Job</i> batch request to the server as described in the general section.                                                                                                                                                                                                                                                                                            | 9 |
| <b>OPTIONS</b>                                                                                                                                                                                                                                                                                                                                                                                        | 9 |
| <i>-h hold_list</i> Defines the types of holds to be placed on the job.                                                                                                                                                                                                                                                                                                                               | 9 |
| The <i>hold_list</i> argument is a string consisting of one or more of the letters "u", "o", or "s" in any combination or the character "n". The hold type associated with each letter is:                                                                                                                                                                                                            | 9 |
| u - <i>USER</i>                                                                                                                                                                                                                                                                                                                                                                                       | 9 |
| o - <i>OTHER</i>                                                                                                                                                                                                                                                                                                                                                                                      | 9 |
| s - <i>SYSTEM</i>                                                                                                                                                                                                                                                                                                                                                                                     | 9 |
| n - <i>None</i>                                                                                                                                                                                                                                                                                                                                                                                       | 9 |
| Repetition of characters is permitted, but "n" may not appear in the same option argument with the other three characters.                                                                                                                                                                                                                                                                            | 9 |
| <b>OPERANDS</b>                                                                                                                                                                                                                                                                                                                                                                                       | 9 |
| The qhold command accepts one or more <i>job_identifier</i> operands of the form:                                                                                                                                                                                                                                                                                                                     | 9 |
| <i>sequence_number</i> [. <i>server_name</i> ][@ <i>server</i> ]                                                                                                                                                                                                                                                                                                                                      |   |
| See the description under "Job Identifier" in section 2.7.6 in this ERS.                                                                                                                                                                                                                                                                                                                              | 9 |
| <b>STANDARD ERROR</b>                                                                                                                                                                                                                                                                                                                                                                                 | 9 |
| The qhold command will write a diagnostic message to standard error for each error occurrence.                                                                                                                                                                                                                                                                                                        | 9 |

|                                                                                                                                      |   |
|--------------------------------------------------------------------------------------------------------------------------------------|---|
| <b>EXIT STATUS</b>                                                                                                                   | 9 |
| Upon successful processing of all the operands presented to the qhold command, the exit status will be a value of zero.              | 9 |
| If the qhold command fails to process any operand, the command exits with a value greater than zero.                                 | 9 |
| <b>SEE ALSO</b>                                                                                                                      | 9 |
| qrls(1B), qalter(1B), qsub(1B), pbs_alterjob(3B), pbs_holdjob(3B), pbs_rlsjob(3B), pbs_job_attributes(7B), pbs_resources_unicos8(7B) | 9 |

|                                                                                                                                                                                                                                                                       |   |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| 5.2.4. Move Job                                                                                                                                                                                                                                                       | 9 |
| <b>NAME</b>                                                                                                                                                                                                                                                           | 9 |
| qmove – move pbs batch job                                                                                                                                                                                                                                            | 9 |
| <b>SYNOPSIS</b>                                                                                                                                                                                                                                                       | 9 |
| qmove destination job_identifier ...                                                                                                                                                                                                                                  | 9 |
| <b>DESCRIPTION</b>                                                                                                                                                                                                                                                    | 9 |
| To move a job is to remove the job from the queue in which it resides and instantiate the job in another queue. The <b>qmove</b> command issues a <i>Move Job</i> batch request to the batch server that currently owns each job specified by <i>job_identifier</i> . | 9 |
| A job in the <b>Running</b> , <b>Transiting</b> , or <b>Exiting</b> state cannot be moved.                                                                                                                                                                            | 9 |
| <b>OPERANDS</b>                                                                                                                                                                                                                                                       | 9 |
| The first operand is the new <i>destination</i> for                                                                                                                                                                                                                   | 9 |
| <i>queue</i>                                                                                                                                                                                                                                                          | 9 |
| <i>@server</i>                                                                                                                                                                                                                                                        | 9 |
| <i>queue@server</i>                                                                                                                                                                                                                                                   | 9 |
| See the PBS ERS section 2.7.3, "Destination Identifiers".                                                                                                                                                                                                             | 9 |
| If the <i>destination</i> operand describes only a queue, then qmove will move jobs into the queue of the specified name at the job's current server.                                                                                                                 | 9 |
| If the <i>destination</i> operand describes only a batch server, then qmove will move jobs into the default queue at that batch server.                                                                                                                               | 9 |
| If the <i>destination</i> operand describes both a queue and a batch server, then qmove will move the jobs into the specified queue at the specified server.                                                                                                          | 9 |
| All following operands are <i>job_identifiers</i> which specify the jobs to be moved to the new <i>destination</i> . The qmove command accepts one or more <i>job_identifier</i> operands of the form:                                                                | 9 |
| <i>sequence_number[.server_name][@server]</i>                                                                                                                                                                                                                         | 9 |
| See the description under "Job Identifier" in section 2.7.6 in this ERS.                                                                                                                                                                                              | 9 |
| <b>STANDARD ERROR</b>                                                                                                                                                                                                                                                 | 9 |
| The qmove command will write a diagnostic messages to standard error for each error occurrence.                                                                                                                                                                       | 9 |
| <b>EXIT STATUS</b>                                                                                                                                                                                                                                                    | 9 |
| Upon successful processing of all the operands presented to the qmove command, the exit status will be a value of zero.                                                                                                                                               | 9 |
| If the qmove command fails to process any operand, the command exits with a value greater than zero.                                                                                                                                                                  | 9 |
| <b>SEE ALSO</b>                                                                                                                                                                                                                                                       | 9 |
| qsub(1B), pbs_movejob(3B)                                                                                                                                                                                                                                             | 9 |

|                                                                                                                                                                                                                                                             |   |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| 5.2.5. Message Job                                                                                                                                                                                                                                          | 9 |
| <b>NAME</b>                                                                                                                                                                                                                                                 | 9 |
| qmsg – send message to pbs batch jobs                                                                                                                                                                                                                       | 9 |
| <b>SYNOPSIS</b>                                                                                                                                                                                                                                             | 9 |
| qmsg [-E] [-O] message_string job_identifier ...                                                                                                                                                                                                            | 9 |
| <b>DESCRIPTION</b>                                                                                                                                                                                                                                          | 9 |
| To send a message to a job is to write a message string into one or more output files of the job. Typically this is done to leave an informative message in the output of the job.                                                                          | 9 |
| The <b>qmsg</b> command writes messages into the files of jobs by sending a <i>Message Job</i> batch request to the batch server that owns the job. The qmsg command does not directly write the message into the files of the job.                         | 9 |
| <b>OPTIONS</b>                                                                                                                                                                                                                                              | 9 |
| -E                Specifies that the message is written to the standard error of each job.                                                                                                                                                                  | 9 |
| -O                Specifies that the message is written to the standard output of each job.                                                                                                                                                                 | 9 |
| If neither the <i>-E</i> nor the <i>-O</i> option is specified, the message will be written to the standard error of the job.                                                                                                                               | 9 |
| <b>OPERANDS</b>                                                                                                                                                                                                                                             | 9 |
| The first operand, <i>message_string</i> , is the message to be written. If the string contains blanks, the string must be quoted. If the final character of the string is not a newline, a newline character will be added when written to the job's file. | 9 |
| All following operands are <i>job_identifiers</i> which specify the jobs to receive the message string. The qmsg command accepts one or more <i>job_identifier</i> operands of the form:                                                                    | 9 |
| <i>sequence_number[.server_name][@server]</i>                                                                                                                                                                                                               | 9 |
| See the description under "Job Identifier" in section 2.7.6 in this ERS.                                                                                                                                                                                    | 9 |
| <b>STANDARD ERROR</b>                                                                                                                                                                                                                                       | 9 |
| The qmsg command will write a diagnostic message to standard error for each error occurrence.                                                                                                                                                               | 9 |
| <b>EXIT STATUS</b>                                                                                                                                                                                                                                          | 9 |
| Upon successful processing of all the operands presented to the qmsg command, the exit status will be a value of zero.                                                                                                                                      | 9 |
| If the qmsg command fails to process any operand, the command exits with a value greater than zero.                                                                                                                                                         | 9 |
| <b>SEE ALSO</b>                                                                                                                                                                                                                                             | 9 |
| qsub(1B), pbs_msgjob(3B)                                                                                                                                                                                                                                    | 9 |

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |   |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| 5.2.6. Order Jobs                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | 9 |
| <b>NAME</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | 9 |
| qorder – exchange order of two pbs batch jobs in a queue.                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | 9 |
| <b>SYNOPSIS</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | 9 |
| qorder job_identifier job_identifier                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | 9 |
| <b>DESCRIPTION</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | 9 |
| To order two jobs is to exchange the jobs positions in the queue or queues in which the jobs resides. The two jobs must be located at the same server. The <b>qorder</b> command issues an <i>Order Jobs</i> batch request to the batch server that currently owns the two jobs specified by the two <i>job_identifier</i> operands. No attribute of the job, such as priority is changed. The impact of interchanging the order with the queue(s) is dependent on local job scheduled policy, contact your systems administrator. | 9 |
| A job in the <b>running</b> state cannot be reordered.                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | 9 |
| <b>OPERANDS</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | 9 |
| Both operands are <i>job_identifiers</i> which specify the jobs to be exchanged. The qorder command accepts two <i>job_identifier</i> operands of the form:<br><i>sequence_number[.server_name][@server]</i>                                                                                                                                                                                                                                                                                                                       | 9 |
| The server specification for the two jobs must agree as to the current location of the two job ids. See the description under "Job Identifier" in section 2.7.6 in this ERS.                                                                                                                                                                                                                                                                                                                                                       | 9 |
| <b>STANDARD ERROR</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | 9 |
| The qorder command will write diagnostic messages to standard error for each error occurrence.                                                                                                                                                                                                                                                                                                                                                                                                                                     | 9 |
| <b>EXIT STATUS</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | 9 |
| Upon successful processing of all the operands presented to the qorder command, the exit status will be a value of zero.                                                                                                                                                                                                                                                                                                                                                                                                           | 9 |
| If the qorder command fails to process any operand, the command exits with a value greater than zero.                                                                                                                                                                                                                                                                                                                                                                                                                              | 9 |
| <b>SEE ALSO</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | 9 |
| qsub(1B), qmove(1B), pbs_orderjob(3B), pbs_movejob(3B)                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | 9 |

|                                                                                                                                                                                              |   |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| 5.2.7. Rerun Job                                                                                                                                                                             | 9 |
| <b>NAME</b>                                                                                                                                                                                  | 9 |
| qrerun – rerun a pbs batch job                                                                                                                                                               | 9 |
| <b>SYNOPSIS</b>                                                                                                                                                                              | 9 |
| qrerun [-W force] job_identifier ...                                                                                                                                                         | 9 |
| <b>DESCRIPTION</b>                                                                                                                                                                           | 9 |
| The <b>qrerun</b> command directs that the specified jobs are to be rerun if possible.                                                                                                       | 9 |
| To rerun a job is to terminate the session leader of the job and return the job to the queued state in the execution queue in which the job currently resides.                               | 9 |
| The qrerun command sends a <i>Rerun Job</i> batch request to the server which owns the job.                                                                                                  | 9 |
| If a job is marked as not rerunable then the rerun request will fail for that job. See the Rerunable attribute and the <i>-r</i> option on the <b>qsub</b> and <b>qalter</b> commands.       | 9 |
| The <i>-W force</i> option, where force is the literal character string <i>force</i> , directs that the job is to be requeued even if the node on which the job is executing is unreachable. | 9 |
| <b>OPERANDS</b>                                                                                                                                                                              | 9 |
| The qrerun command accepts one or more <i>job_identifier</i> operands of the form:                                                                                                           | 9 |
| <i>sequence_number[.server_name][@server]</i>                                                                                                                                                | 9 |
| See the description under "Job Identifier" in section 2.7.6 in this ERS.                                                                                                                     | 9 |
| <b>STANDARD ERROR</b>                                                                                                                                                                        | 9 |
| The qrerun command will write a diagnostic message to standard error for each error occurrence.                                                                                              | 9 |
| <b>EXIT STATUS</b>                                                                                                                                                                           | 9 |
| Upon successful processing of all the operands presented to the qrerun command, the exit status will be a value of zero.                                                                     | 9 |
| If the qrerun command fails to process any operand, the command exits with a value greater than zero.                                                                                        | 9 |
| <b>SEE ALSO</b>                                                                                                                                                                              | 9 |
| qsub(1B), qalter(1B), pbs_alterjob(3B), pbs_rerunjob(3B)                                                                                                                                     | 9 |

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |                                                                                                                                                                                                                                                                         |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 5.2.8. Release Job                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | 9                                                                                                                                                                                                                                                                       |
| <b>NAME</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | 9                                                                                                                                                                                                                                                                       |
| qrls – release hold on pbs batch jobs                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | 9                                                                                                                                                                                                                                                                       |
| <b>SYNOPSIS</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | 9                                                                                                                                                                                                                                                                       |
| qrls [-h hold_list] job_identifier ...                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | 9                                                                                                                                                                                                                                                                       |
| <b>DESCRIPTION</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | 9                                                                                                                                                                                                                                                                       |
| The <b>qrls</b> command removes or releases holds which exist on batch jobs.                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | 9                                                                                                                                                                                                                                                                       |
| A job may have one or more types of holds which make the job ineligible for execution. The types of holds are <i>USER</i> , <i>OTHER</i> , and <i>SYSTEM</i> . The different types of holds may require that the user issuing the qrls command have special privilege. Typically, the owner of the job will be able to remove a <i>USER</i> hold, but not an <i>OTHER</i> or <i>SYSTEM</i> hold. An Attempt to release a hold for which the user does not have the correct privilege is an error and no holds will be released for that job. | 9                                                                                                                                                                                                                                                                       |
| If no <i>-h</i> option is specified, the <i>USER</i> hold will be released.                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | 9                                                                                                                                                                                                                                                                       |
| If the job has no execution_time pending, the job will change to the <b>queued</b> state. If an execution_time is still pending, the job will change to the <b>waiting</b> state.                                                                                                                                                                                                                                                                                                                                                            | 9                                                                                                                                                                                                                                                                       |
| If the sched_hint attribute is set, when the job is returned to <b>queued</b> state, it may be given preference in selection for execution depending on site policy.                                                                                                                                                                                                                                                                                                                                                                         | 9                                                                                                                                                                                                                                                                       |
| The qrls command sends a <i>Release Job</i> batch request to the server which owns the job.                                                                                                                                                                                                                                                                                                                                                                                                                                                  | 9                                                                                                                                                                                                                                                                       |
| <b>OPTIONS</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | 9                                                                                                                                                                                                                                                                       |
| -h hold_list                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | Defines the types of hold to be released from the jobs. The <i>hold_list</i> option argument is a string consisting of one or more of the letters "u", "o", an "s" in any combination, or one or more of the letters "n". The hold type associated with each letter is: |
| u – <i>USER</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | 9                                                                                                                                                                                                                                                                       |
| o – <i>OTHER</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | 9                                                                                                                                                                                                                                                                       |
| s – <i>SYSTEM</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | 9                                                                                                                                                                                                                                                                       |
| n – <i>None</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | 9                                                                                                                                                                                                                                                                       |
|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Repetition of characters is permitted, but "n" may not appear in the same option argument with the other three characters.                                                                                                                                              |
|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | 9                                                                                                                                                                                                                                                                       |
| <b>OPERANDS</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | 9                                                                                                                                                                                                                                                                       |
| The qrls command accepts one or more <i>job_identifier</i> operands of the form:                                                                                                                                                                                                                                                                                                                                                                                                                                                             | 9                                                                                                                                                                                                                                                                       |
| <i>sequence_number</i> [@server_name][@server]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | 9                                                                                                                                                                                                                                                                       |
| See the description under "Job Identifier" in section 2.7.6 in this ERS.                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | 9                                                                                                                                                                                                                                                                       |
| <b>STANDARD ERROR</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | 9                                                                                                                                                                                                                                                                       |
| The qrls command will write a diagnostic message to standard error for each error occurrence.                                                                                                                                                                                                                                                                                                                                                                                                                                                | 9                                                                                                                                                                                                                                                                       |
| <b>EXIT STATUS</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | 9                                                                                                                                                                                                                                                                       |
| Upon successful processing of all the operands presented to the qrls command, the exit status will be a value of zero.                                                                                                                                                                                                                                                                                                                                                                                                                       | 9                                                                                                                                                                                                                                                                       |
| If the qrls command fails to process any operand, the command exits with a value greater than zero.                                                                                                                                                                                                                                                                                                                                                                                                                                          | 9                                                                                                                                                                                                                                                                       |
| <b>SEE ALSO</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | 9                                                                                                                                                                                                                                                                       |

qsub(1B), qalter(1B), qhold(1B), pbs\_alterjob(3B), pbs\_holdjob(3B), and pbs\_rlsjob(3B).

9



## 5.2.9. Select Jobs 9

**NAME** 9

qselect – select pbs batch jobs 9

**SYNOPSIS** 9

qselect [-a [op]date\_time] [-A account\_string] [-c [op]interval] [-h hold\_list] [-l resource\_list] [-N name] [-p [op]priority] [-q destination] [-r rerun] [-s states] [-u user\_list] 9

**DESCRIPTION** 9

The **qselect** command provides a method to list the job identifier of those jobs which meet a list of selection criteria. The selection is accomplished by sending a *Select Jobs* batch request to the default server or the server specified by the *-q* option. Jobs are selected from those owned by a single server. 9

When qselect successfully completes, it will have written to standard output a list of zero or more jobs which meet the criteria specified by the options. Each option acts as a filter restricting the number of jobs which might be listed. With no options, the qselect command will list all jobs at the server which the user is authorized to list (query status of). The *-u* option may be used to limit the selection to jobs owned by this user or other specified users. Operators and system administrators have the privilege to select all jobs. Other uses access depends on the setting of the server attribute *query\_other\_jobs*. 9

**OPTIONS** 9

When an option is specified with a optional *op* component to the option argument, then *op* specifies a relation between the value of a certain job attribute and the value component of the option argument. If an *op* is allowable on an option, then the description of the option letter will indicate the *op* is allowable. The only acceptable strings for the *op* component, and the relation the string indicates, are shown in the following list: 9

- .eq. the value represented by the attribute of the job is equal to the value represented by the option argument. 9
- .ne. the value represented by the attribute of the job is not equal to the value represented by the option argument. 9
- .ge. the value represented by the attribute of the job is greater than or equal to the value represented by the option argument. 9
- .gt. the value represented by the attribute of the job is greater than the value represented by the option argument. 9
- .le. the value represented by the attribute of the job is less than or equal to the value represented by the option argument. 9
- .lt. the value represented by the attribute of the job is less than the value represented by the option argument. 9

*-a* [op]date\_time 9

Restricts selection to a specific time, or a range of times. 9

The qselect command selects only jobs for which the value of the *Execution\_Time* attribute is related to the *date\_time* argument by the optional *op* operator. The *date\_time* argument is in the form of the *date\_time* operand of the **touch**(1) command: *[[CC]YY]MMDDhhmm[.SS]* 9

where the MM is the two digits for the month, DD is the day of the month, hh is the hour, mm is the minute, and the optional SS is the seconds. CC is the century and YY the year. 9

If *op* is not specified, jobs will be selected for which the *Execution\_Time* and *date\_time* values are equal. If *op* is specified, jobs will be selected according to 9

|                                                                                                                                                                                                                                                                                                              |   |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| the following definitions:                                                                                                                                                                                                                                                                                   | 9 |
| .eq. Execution_Time attribute is equal to the <i>date_time</i> argument.                                                                                                                                                                                                                                     | 9 |
| .ne. Execution_Time attribute is not equal to the <i>date_time</i> argument.                                                                                                                                                                                                                                 | 9 |
| .ge. execution_Time attribute is greater than (after) or equal to the <i>date_time</i> argument.                                                                                                                                                                                                             | 9 |
| .gt. Execution_Time attribute is greater than (after) the <i>date_time</i> argument.                                                                                                                                                                                                                         | 9 |
| .le. Execution_Time attribute is less than (before) or equal to the <i>date_time</i> argument.                                                                                                                                                                                                               | 9 |
| .lt. Execution_Time attribute is less than (before) the <i>date_time</i> argument.                                                                                                                                                                                                                           | 9 |
| -A <i>account_string</i>                                                                                                                                                                                                                                                                                     | 9 |
| Restricts selection to jobs whose Account_Name attribute matches the specified <i>account_string</i> .                                                                                                                                                                                                       | 9 |
| -c [ <i>op</i> ] <i>interval</i>                                                                                                                                                                                                                                                                             | 9 |
| Restricts selection to jobs whose Checkpoint interval attribute matches the specified relationship.                                                                                                                                                                                                          | 9 |
| The values of the Checkpoint attribute are defined to have the following ordered relationship:                                                                                                                                                                                                               | 9 |
| n > s > c=minutes > c > u                                                                                                                                                                                                                                                                                    | 9 |
| If the optional <i>op</i> is not specified, jobs will be selected whose Checkpoint attribute is equal to the <i>interval</i> argument. If <i>op</i> is specified, jobs will be selected according to:                                                                                                        | 9 |
| .eq. Checkpoint attribute of the job is equal to the <i>interval</i> argument.                                                                                                                                                                                                                               | 9 |
| .ne. Checkpoint attribute of the job is not equal to the <i>interval</i> argument.                                                                                                                                                                                                                           | 9 |
| .ge. Checkpoint attribute of the job is greater than or equal to the <i>interval</i> argument.                                                                                                                                                                                                               | 9 |
| .gt. Checkpoint attribute of the job is greater than the <i>interval</i> argument.                                                                                                                                                                                                                           | 9 |
| .le. Checkpoint attribute of the job is less than or equal to the <i>interval</i> argument.                                                                                                                                                                                                                  | 9 |
| .lt. Checkpoint attribute of the job is less than the <i>interval</i> argument.                                                                                                                                                                                                                              | 9 |
| For an interval value of "u", only ".eq." and ".ne." are valid.                                                                                                                                                                                                                                              | 9 |
| -h <i>hold_list</i>                                                                                                                                                                                                                                                                                          | 9 |
| Restricts the selection of jobs to those with a specific set of hold types. Only those jobs will be selected whose Hold_Types attribute exactly match the value of the <i>hold_list</i> argument.                                                                                                            | 9 |
| The <i>hold_list</i> argument is a string consisting of one or more occurrences the single letter <i>n</i> , or one or more of the letters <i>u</i> , <i>o</i> , or <i>s</i> in any combination. If letters are duplicated, they are treated as if they occurred once. The letters represent the hold types: | 9 |
| n – none                                                                                                                                                                                                                                                                                                     | 9 |
| u – user                                                                                                                                                                                                                                                                                                     | 9 |
| o – other                                                                                                                                                                                                                                                                                                    | 9 |
| s – system                                                                                                                                                                                                                                                                                                   | 9 |
| -l <i>resource_list</i>                                                                                                                                                                                                                                                                                      | 9 |
| Restricts selection of jobs to those with specified resource amounts.                                                                                                                                                                                                                                        | 9 |
| Only those jobs will be selected whose Resource_List attribute matches the specified relation with each resource and value listed in the <i>resource_list</i> argument. The <i>resource_list</i> is in the following format:                                                                                 | 9 |
| <i>resource_name</i> <i>op</i> <i>value</i> [ , <i>resource_name</i> <i>op</i> <i>val</i> , . . . ]                                                                                                                                                                                                          | 9 |
| The relation operator <i>op</i> must be present.                                                                                                                                                                                                                                                             | 9 |

|                 |                                                                                                                                                                                                                            |   |
|-----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
|                 | When comparing the values of resources, the following definitions for the operator apply:                                                                                                                                  | 9 |
|                 | .eq. the resource value in the Resource_List attribute of the job equals the value specified in <i>resource_list</i> .                                                                                                     | 9 |
|                 | .ne. the resource value in the Resource_List attribute of the job is not equal to the value specified in <i>resource_list</i> .                                                                                            | 9 |
|                 | .ge. the resource value in the Resource_List attribute of the job is greater than or equal to the value specified in <i>resource_list</i> .                                                                                | 9 |
|                 | .gt. the resource value in the Resource_List attribute of the job is greater than the value specified in <i>resource_list</i> .                                                                                            | 9 |
|                 | .le. the resource value in the Resource_List attribute of the job is less than or equal to the value specified in <i>resource_list</i> .                                                                                   | 9 |
|                 | .lt. the resource value in the Resource_List attribute of the job is less than the value specified in <i>resource_list</i> .                                                                                               | 9 |
| -N name         | Restricts selection of jobs to those with a specific name.                                                                                                                                                                 | 9 |
| -p [op]priority | Restricts selection of jobs to those with a priority that matches the specified relationship. If <i>op</i> is not specified, jobs are selected for which the job Priority attribute is equal to the <i>priority</i>        | 9 |
|                 | If the <i>op</i> is specified, the relationship is defined as:                                                                                                                                                             | 9 |
|                 | .eq. Priority attribute is equal to the value of the <i>priority</i> argument.                                                                                                                                             | 9 |
|                 | .ne. Priority attribute is not equal to the value of the <i>priority</i> argument.                                                                                                                                         | 9 |
|                 | .ge. Priority attribute is greater than or equal to the value of the <i>priority</i> argument.                                                                                                                             | 9 |
|                 | .gt. Priority attribute is greater than the value of the <i>priority</i> argument.                                                                                                                                         | 9 |
|                 | .le. Priority attribute is less than or equal to the value of the <i>priority</i> argument.                                                                                                                                | 9 |
|                 | .lt. Priority attribute is less than the value of the <i>priority</i> argument.                                                                                                                                            | 9 |
| -q destination  | Restricts selection to those jobs residing at the specified destination.                                                                                                                                                   | 9 |
|                 | The <i>destination</i> may be of one of the following three forms:                                                                                                                                                         | 9 |
|                 | <i>queue</i>                                                                                                                                                                                                               | 9 |
|                 | @ <i>server</i>                                                                                                                                                                                                            | 9 |
|                 | <i>queue@server</i>                                                                                                                                                                                                        | 9 |
|                 | If the <i>-q</i> option is not specified, jobs will be selected from the default server. See the ERS section 2.7.4 for a definition of the default server.                                                                 | 9 |
|                 | If the <i>destination</i> describes only a queue, only jobs in that queue on the default batch server will be selected.                                                                                                    | 9 |
|                 | If the <i>destination</i> describes only a server, then jobs in all queues on that server will be selected.                                                                                                                | 9 |
|                 | If the <i>destination</i> describes both a queue and a server, then only jobs in the named queue on the named server will be selected.                                                                                     | 9 |
| -r rerun        | Restricts selection of jobs to those with the specified Rerunable attribute. The option argument must be a single character. The following two characters are supported by PBS: <i>y</i> and <i>n</i> .                    | 9 |
| -s states       | Restricts job selection to those in the specified states.                                                                                                                                                                  | 9 |
|                 | The <i>states</i> argument is a character string which consists of any combination of the characters: <i>E</i> , <i>H</i> , <i>Q</i> , <i>R</i> , <i>T</i> , and <i>W</i> . [A repeated character will be accepted, but no | 9 |

|                                                                                                                                                                                                                                                                                                                                                                                                                                                               |   |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| additional meaning is assigned to it.]                                                                                                                                                                                                                                                                                                                                                                                                                        | 9 |
| The characters in the <i>states</i> argument have the following interpretation:                                                                                                                                                                                                                                                                                                                                                                               | 9 |
| E the Exiting state.                                                                                                                                                                                                                                                                                                                                                                                                                                          | 9 |
| H the Held state.                                                                                                                                                                                                                                                                                                                                                                                                                                             | 9 |
| Q the Queued state.                                                                                                                                                                                                                                                                                                                                                                                                                                           | 9 |
| R the Running state.                                                                                                                                                                                                                                                                                                                                                                                                                                          | 9 |
| T the Transiting state.                                                                                                                                                                                                                                                                                                                                                                                                                                       | 9 |
| W the Waiting state.                                                                                                                                                                                                                                                                                                                                                                                                                                          | 9 |
| Jobs will be selected which are in any of the specified states.                                                                                                                                                                                                                                                                                                                                                                                               | 9 |
| -u <i>user_list</i> Restricts selection to jobs owned by the specified user names.                                                                                                                                                                                                                                                                                                                                                                            | 9 |
| This provides a means of limiting the selection to jobs owned by one or more users. The ability to select jobs owned by others is controllable by the server attribute <i>query_other_jobs</i> . Mapping between user names on different hosts and validation of privilege to access the specified user name is discussed under the server. This option may also be used by batch operators and batch administrators to select jobs belonging to other users. | 9 |
| The syntax of the <i>user_list</i> is:                                                                                                                                                                                                                                                                                                                                                                                                                        | 9 |
| <i>user_name</i> [@ <i>host</i> ][, <i>user_name</i> [@ <i>host</i> ],...]                                                                                                                                                                                                                                                                                                                                                                                    | 9 |
| Host names may be wild carded on the left end, e.g. "*.nasa.gov". <i>User_name</i> without a "@host" is equivalent to " <i>user_name</i> @*", that is at any host. Jobs will be selected which are owned by the listed users at the corresponding hosts.                                                                                                                                                                                                      | 9 |
| <b>STANDARD OUTPUT</b>                                                                                                                                                                                                                                                                                                                                                                                                                                        | 9 |
| The list of job identifiers of selected jobs is written to standard output. Each job identifier is separated by white space. Each job identifier is of the form:                                                                                                                                                                                                                                                                                              | 9 |
| <i>sequence_number.server_name@server</i>                                                                                                                                                                                                                                                                                                                                                                                                                     | 9 |
| Where <i>sequence_number.server</i> is the identifier assigned at submission time, see <b>qsub</b> . <i>@server</i> identifies the server which currently owns the job.                                                                                                                                                                                                                                                                                       | 9 |
| <b>STANDARD ERROR</b>                                                                                                                                                                                                                                                                                                                                                                                                                                         | 9 |
| The <i>qselect</i> command will write a diagnostic message to standard error for each error occurrence.                                                                                                                                                                                                                                                                                                                                                       | 9 |
| <b>EXIT STATUS</b>                                                                                                                                                                                                                                                                                                                                                                                                                                            | 9 |
| Upon successful processing of all options presented to the <i>qselect</i> command, the exit status will be a value of zero.                                                                                                                                                                                                                                                                                                                                   | 9 |
| If the <i>qselect</i> command fails to process any option, the command exits with a value greater than zero.                                                                                                                                                                                                                                                                                                                                                  | 9 |
| <b>SEE ALSO</b>                                                                                                                                                                                                                                                                                                                                                                                                                                               | 9 |
| <i>qsub</i> (1B), <i>qstat</i> (1B), <i>pbs_selectjob</i> (3B), <i>pbs_selstat</i> (3B), <i>pbs_statjob</i> (3B)                                                                                                                                                                                                                                                                                                                                              | 9 |

## 5.2.10. Signal Job 9

**NAME** 9

qsig – signal pbs batch job 9

**SYNOPSIS** 9

qsig [-s signal] job\_identifier ... 9

**DESCRIPTION** 9

The **qsig** command requests that a signal be sent to executing batch jobs. The signal is sent to the session leader of the job. 9

If the `-s` option is not specified, '**SIGTERM**' is sent. The request to signal a batch job will be rejected if: 9

- The user is not authorized to signal the job. 9
- The job is not in the **running** state. 9
- The requested signal is not supported by the system upon which the job is executing. 9

The **qsig** command sends a *Signal Job* batch request to the server which owns the job. 9

**OPTIONS** 9

`-s signal` Declares which signal is sent to the job. 9

The *signal* argument is either a signal name, e.g. **SIGKILL**, the signal name without the **SIG** prefix, e.g. **KILL**, or a unsigned signal number, e.g. **9**. The signal name **SIGNULL** is allowed; the server will send the signal 0 to the job which will have no effect. Not all signal names will be recognized by **qsig**. If it doesn't recognize the signal name, try issuing the signal number instead. 9

Two special signal names, "suspend" and "resume", [note, all lower case], are used to suspend and resume jobs. When suspended, a job continues to occupy system resources but is not executing and is not charged for walltime. Manager or operator privilege is required to suspend or resume a job. For systems other than Unicos, the SIGSTOP and SIGCONT signals are used to suspend and resume the job. 9

If the server receives a *Signal Job* batch request with a signal that is unsupported on the server host, the server will reject the request. 9

**OPERANDS** 9

The **qsig** command accepts one or more *job\_identifier* operands of the form: 9

*sequence\_number*[.*server\_name*][@*server*]

See the description under "Job Identifier" in section 2.7.6 in this ERS. 9

**STANDARD ERROR** 9

The **qsig** command will write a diagnostic messages to standard error for each error occurrence. 9

**EXIT STATUS** 9

Upon successful processing of all the operands presented to the **qsig** command, the exit status will be a value of zero. 9

If the **qsig** command fails to process any operand, the command exits with a value greater than zero. 9

**SEE ALSO**

qsub(1B), pbs\_sigjob(3B), pbs\_resources\_\*(7B) where \* is system type, and the PBS ERS. 9

|                                                                                                                                                                                                                                                                                                         |   |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| 5.2.11. Status Jobs, Queues, or Server                                                                                                                                                                                                                                                                  | 9 |
| <b>NAME</b>                                                                                                                                                                                                                                                                                             | 9 |
| qstat – show status of pbs batch jobs                                                                                                                                                                                                                                                                   | 9 |
| <b>SYNOPSIS</b>                                                                                                                                                                                                                                                                                         | 9 |
| qstat [-f][-W site_specific] [job_identifier...   destination...]                                                                                                                                                                                                                                       | 9 |
| qstat [-a   -i   -r] [-n] [-s] [-G   -M] [-R] [-u user_list] [job_identifier...   destination...]                                                                                                                                                                                                       | 9 |
| qstat -Q [-f][-W site_specific] [destination...]                                                                                                                                                                                                                                                        | 9 |
| qstat -q [-G   -M] [destination...]                                                                                                                                                                                                                                                                     | 9 |
| qstat -B [-f][-W site_specific] [server_name...]                                                                                                                                                                                                                                                        | 9 |
| <b>DESCRIPTION</b>                                                                                                                                                                                                                                                                                      | 9 |
| The <b>qstat</b> command is used to request the status of jobs, queues, or a batch server. The requested status is written to standard out. The status is obtained by sending a <i>Job Status</i> , a <i>Queue Status</i> , or a <i>Server Status</i> batch request to the appropriate server.          | 9 |
| When requesting job status, synopsis format 1 or 2, <b>qstat</b> will output information about each <i>job_identifier</i> or all jobs at each <i>destination</i> . Jobs for which the user does not have status privilege are not displayed.                                                            | 9 |
| When requesting queue or server status, synopsis format 3 through 5, qstat will output information about each <i>destination</i> . The syntax using the [-a   -i   -r] line or the -q line are extensions to POSIX.                                                                                     | 9 |
| <b>OPTIONS</b>                                                                                                                                                                                                                                                                                          | 9 |
| -f Specifies that a full status display be written to standard out.                                                                                                                                                                                                                                     | 9 |
| -a "All" jobs are displayed in the alternative format, see the Standard Output section. If the operand is a destination id, all jobs at that destination are displayed. If the operand is a job id, information about that job is displayed.                                                            | 9 |
| -i Job status is displayed in the alternative format. For a destination id operand, status for jobs at that destination which are not running are displayed. This includes jobs which are queued, held or waiting. If an operand is a job id, status for that job is displayed regardless of its state. | 9 |
| -r If an operand is a job id, status for that job is displayed. For a destination id operand, status for jobs at that destination which are running are displayed, this includes jobs which are suspended. If an operand is a job id, status for that job is displayed.                                 | 9 |
| -n In addition to the basic information, nodes allocated to a job are listed.                                                                                                                                                                                                                           | 9 |
| -s In addition to the basic information, any comment provided by the batch administrator or scheduler is shown.                                                                                                                                                                                         | 9 |
| -G Show size information in giga-bytes.                                                                                                                                                                                                                                                                 | 9 |
| -M Show size information, disk or memory in mega-words. A word is considered to be 8 bytes.                                                                                                                                                                                                             | 9 |
| -R In addition to other information, disk reservation information is shown. Not applicable to all systems.                                                                                                                                                                                              | 9 |
| -u Job status is displayed in the alternative format. If an operand is a job id, status for that job is displayed. For a destination id operand, status for jobs at that destination which are owned by the user(s) listed in <i>user_list</i> are displayed. The syn-                                  | 9 |

- tax of the *user\_list* is: 9
- user\_name[@host][,user\_name[@host],...]* 9
- Host names may be wild carded on the left end, e.g. *\*.nasa.gov*. *User\_name* 9
- without a *"@host"* is equivalent to *"user\_name@\*"*, that is at any host. 9
- Q Specifies that the request is for queue status and that the operands are destina- 9
- tion identifiers. 9
- q Specifies that the request is for queue status which should be shown in the alter- 9
- native format. 9
- B Specifies that the request is for batch server status and that the operands are the 9
- names of servers. 9

**OPERANDS** 9

If neither the *-Q* nor the *-B* option is given, the operands on the *qstat* command must be either job identifiers or destinations identifiers. 9

If the operand is a job identifier, it must be in the following form: 9

*sequence\_number[.server\_name][@server]* 9

where *sequence\_number.server\_name* is the job identifier assigned at submittal time, see **qsub**. If the *.server\_name* is omitted, the name of the default server will be used. If *@server* is supplied, the request will be for the job identifier currently at that Server. See ERS sections 2.7.6 and 2.7.3 for more details on job identifiers and batch destinations. 9

If the operand is a destination identifier, it is one of the following three forms: 9

*queue* 9

*@server* 9

*queue@server* 9

If *queue* is specified, the request is for status of all jobs in that queue at the default server. If the *@server* form is given, the request is for status of all jobs at that server. If a full destination identifier, *queue@server*, is given, the request is for status of all jobs in the named queue at the named server. 9

If the *-Q* option is given, the operands are destination identifiers as specified above. If *queue* is specified, the status of that queue at the default server will be given. If *queue@server* is specified, the status of the named queue at the named server will be given. If *@server* is specified, the status of all queues at the named server will be given. If no destination is specified, the status of all queues at the default server will be given. 9

If the *-B* option is given, the operand is the name of a server. 9

**STANDARD OUTPUT** 9

Displaying Job Status 9

If job status is being displayed in the default format and the *-f* option is not specified, the following items are displayed on a single line, in the specified order, separated by white space: 9

- the job identifier assigned by PBS. 9
- the job name given by the submitter. 9
- the job owner 9
- the CPU time used 9
- the job state: 9
  - E - Job is exiting after having run. 9
  - H - Job is held. 9
  - Q - job is queued, eligible to run or routed. 9
  - R - job is running. 9
  - T - job is being moved to new location. 9
  - W - job is waiting for its execution time 9



|                                                                                                                                                                                                                                                                                                                                                |   |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| (-a option) to be reached.                                                                                                                                                                                                                                                                                                                     | 9 |
| S - (Unicos only) job is suspend.                                                                                                                                                                                                                                                                                                              | 9 |
| - the queue in which the job resides                                                                                                                                                                                                                                                                                                           | 9 |
| If job status is being displayed and the <i>-f</i> option is specified, the output will depend on whether <b>qstat</b> was compiled to use a <b>Tcl</b> interpreter. See the configuration section for details. If <b>Tcl</b> is not being used, full display for each job consists of the header line:                                        | 9 |
| <i>Job Id:</i> job identifier                                                                                                                                                                                                                                                                                                                  | 9 |
| Followed by one line per job attribute of the form:                                                                                                                                                                                                                                                                                            | 9 |
| <i>attribute_name = value</i>                                                                                                                                                                                                                                                                                                                  | 9 |
| The attribute name is indented 4 spaces. There is a single space on each side of the equal sign. Long values wrap either at column 78 or following a comma beyond which the next comma separated segment will not fit before column 79. Continuation lines are indented by a tab (8 spaces). There is blank line following the last attribute. | 9 |
| If any of the options -a, -i, -r, -u, -n, -s, -G or -M are provided, the alternative display format for jobs is used. The following items are displayed on a single line, in the specified order, separated by white space:                                                                                                                    | 9 |
| - the job identifier assigned by PBS.                                                                                                                                                                                                                                                                                                          | 9 |
| - the job owner.                                                                                                                                                                                                                                                                                                                               | 9 |
| - The queue in which the job currently resides.                                                                                                                                                                                                                                                                                                | 9 |
| - The job name given by the submitter.                                                                                                                                                                                                                                                                                                         | 9 |
| - The session id (if the job is running).                                                                                                                                                                                                                                                                                                      | 9 |
| - The number of nodes requested by the job.                                                                                                                                                                                                                                                                                                    | 9 |
| - The number of cpus or tasks requested by the job.                                                                                                                                                                                                                                                                                            | 9 |
| - The amount of memory requested by the job.                                                                                                                                                                                                                                                                                                   | 9 |
| - Either the cpu time, if specified, or wall time requested by the job, (hh:mm).                                                                                                                                                                                                                                                               | 9 |
| - The job's current state.                                                                                                                                                                                                                                                                                                                     | 9 |
| - The amount of cpu time or wall time used by the job (hh:mm).                                                                                                                                                                                                                                                                                 | 9 |
| If the -R option is provided, the line contains:                                                                                                                                                                                                                                                                                               | 9 |
| - the job identifier assigned by PBS.                                                                                                                                                                                                                                                                                                          | 9 |
| - the job owner.                                                                                                                                                                                                                                                                                                                               | 9 |
| - The queue in which the job currently resides.                                                                                                                                                                                                                                                                                                | 9 |
| - The number of nodes requested by the job.                                                                                                                                                                                                                                                                                                    | 9 |
| - The number of cpus or tasks requested by the job.                                                                                                                                                                                                                                                                                            | 9 |
| - The amount of memory requested by the job.                                                                                                                                                                                                                                                                                                   | 9 |
| - Either the cpu time or wall time requested by the job.                                                                                                                                                                                                                                                                                       | 9 |
| - The job's current state.                                                                                                                                                                                                                                                                                                                     | 9 |
| - The amount of cpu time or wall time used by the job.                                                                                                                                                                                                                                                                                         | 9 |
| - The amount of SRFS space requested on the big file system.                                                                                                                                                                                                                                                                                   | 9 |
| - The amount of SRFS space requested on the fast file system.                                                                                                                                                                                                                                                                                  | 9 |
| - The amount of space requested on the parallel I/O file system.                                                                                                                                                                                                                                                                               | 9 |
| The last three fields may not contain useful information at all sites or on all systems.                                                                                                                                                                                                                                                       | 9 |

### Displaying Queue Status 9

If queue status is being displayed and the *-f* option was not specified, the following items are displayed on a single line, in the specified order, separated by white space: 9

|                                                                                                                                                                                                                                                                                                                 |   |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| - the queue name                                                                                                                                                                                                                                                                                                | 9 |
| - the maximum number of jobs that may be run in the queue concurrently                                                                                                                                                                                                                                          | 9 |
| - the total number of jobs in the queue                                                                                                                                                                                                                                                                         | 9 |
| - the enable or disabled status of the queue                                                                                                                                                                                                                                                                    | 9 |
| - the started or stopped status of the queue                                                                                                                                                                                                                                                                    | 9 |
| - for each job state, the name of the state and the number of jobs in the queue in that state.                                                                                                                                                                                                                  | 9 |
| - the type of queue, execution or routing.                                                                                                                                                                                                                                                                      | 9 |
| If queue status is being displayed and the <i>-f</i> option is specified, the output will depend on whether <b>qstat</b> was compiled to use a <b>Tcl</b> interpreter. See the configuration section for details. If <b>Tcl</b> is not being used, the full display for each queue consists of the header line: | 9 |
| <i>Queue:</i> queue_name                                                                                                                                                                                                                                                                                        | 9 |
| Followed by one line per queue attribute of the form:                                                                                                                                                                                                                                                           | 9 |
| <i>attribute_name = value</i>                                                                                                                                                                                                                                                                                   | 9 |
| The queue attributes are listed in the same format as job attributes.                                                                                                                                                                                                                                           | 9 |
| If the <i>-q</i> option is specified, queue information is displayed in the alternative format: The following information is displayed on a single line:                                                                                                                                                        | 9 |
| - the queue name                                                                                                                                                                                                                                                                                                | 9 |
| - the maximum amount of memory a job in the queue may request                                                                                                                                                                                                                                                   | 9 |
| - the maximum amount of cpu time a job in the queue may request                                                                                                                                                                                                                                                 | 9 |
| - the maximum amount of wall time a job in the queue may request                                                                                                                                                                                                                                                | 9 |
| - the maximum amount of nodes a job in the queue may request                                                                                                                                                                                                                                                    | 9 |
| - the number of jobs in the queue in the running state                                                                                                                                                                                                                                                          | 9 |
| - the number of jobs in the queue in the queued state                                                                                                                                                                                                                                                           | 9 |
| - the maximum number (limit) of jobs that may be run in the queue concurrently                                                                                                                                                                                                                                  | 9 |
| - the state of the queue given by a pair of letters:                                                                                                                                                                                                                                                            | 9 |
| - either the letter E if the queue is Enabled or D if Disabled, and                                                                                                                                                                                                                                             | 9 |
| - either the letter R if the queue is Running (started) or S if Stopped.                                                                                                                                                                                                                                        | 9 |
| <br>Displaying Server Status                                                                                                                                                                                                                                                                                    | 9 |
| If batch server status is being displayed and the <i>-f</i> option is not specified, the following items are displayed on a single line, in the specified order, separated by white space:                                                                                                                      | 9 |
| - the server name                                                                                                                                                                                                                                                                                               | 9 |
| - the maximum number of jobs that the server may run concurrently                                                                                                                                                                                                                                               | 9 |
| - the total number of jobs currently managed by the server                                                                                                                                                                                                                                                      | 9 |
| - the status of the server                                                                                                                                                                                                                                                                                      | 9 |
| - for each job state, the name of the state and the number of jobs in the server in that state                                                                                                                                                                                                                  | 9 |
| If server status is being displayed and the <i>-f</i> option is specified, the output will depend on whether <b>qstat</b> was compiled to use a <b>Tcl</b> interpreter. See the configuration section for details. If <b>Tcl</b> is not being used, the full display for the server consist of the header line: | 9 |
| <i>Server:</i> server name                                                                                                                                                                                                                                                                                      | 9 |
| Followed by one line per server attribute of the form:                                                                                                                                                                                                                                                          | 9 |
| <i>attribute_name = value</i>                                                                                                                                                                                                                                                                                   | 9 |
| The server attributes are listed in the same format as job attributes.                                                                                                                                                                                                                                          | 9 |

**STANDARD ERROR** 9  
 The `qstat` command will write a diagnostic message to standard error for each error occurrence. 9

**CONFIGURATION** 9

If `qstat` is compiled with an option to include a **Tcl** interpreter, using the `-f` flag to get a full display causes a check to be made for a script file to use to output the requested information. 9  
 The first location checked is `$HOME/.qstatrc`. If this does not exist, the next location checked is administrator configured. If one of these is found, a **Tcl** interpreter is started and the script file is passed to it along with three global variables. The command line arguments are split into two variable named **flags** and **operands**. The status information is passed in a variable named **objects**. All of these variables are **Tcl** lists. The **flags** list contains the name of the command (usually "qstat") as its first element. Any other elements are command line option flags with any options they use, presented in the order given on the command line. They are broken up individually so that if two flags are given together on the command line, they are separated in the list. For example, if the user typed 9

`qstat -QfWbigdisplay` 9

the **flags** list would contain 9

`qstat -Q -f -W bigdisplay` 9

The **operands** list contains all other command line arguments following the flags. There will always be at least one element in **operands** because if no operands are typed by the user, the default destination or server name is used. The **objects** list contains all the information retrieved from the server(s) so the Tcl interpreter can run once to format the entire output. This list has the same number of elements as the **operands** list. Each element is another list with two elements. The first element is a string giving the type of objects to be found in the second. The string can take the values "server", "queue", "job" or "error". The second element will be a list in which each element is a single batch status object of the type given by the string discussed above. In the case of "error", the list will be empty. Each object is again a list. The first element is the name of the object. The second is a list of attributes. The third element will be the object text. All three of these object elements correspond with fields in the structure `batch_status` which is described in detail for each type of object by the man pages for **pbs\_statjob(3)**, **pbs\_statque(3)**. Each attribute in the second element list whose elements correspond with the `attr1` structure. Each will be a list with two elements. The first will be the attribute name and the second will be the attribute value. 9

**EXIT STATUS** 9

Upon successful processing of all the operands presented to the `qstat` command, the exit status will be a value of zero. 9

If the `qstat` command fails to process any operand, the command exits with a value greater than zero. 9

**SEE ALSO** 9

`qalter(1B)`, `qsub(1B)`, `pbs_alterjob(3B)`, `pbs_statjob(3B)`, `pbs_statque(3B)`, `pbs_statserver(3B)`, `pbs_submit(3B)`, `pbs_job_attributes(7B)`, `pbs_queue_attributes(7B)`, `pbs_server_attributes(7B)`, `pbs_resources_*(7B)` where \* is system type, and the PBS ERS. 9

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |   |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| 5.2.12. Submit Job                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | 9 |
| <b>NAME</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | 9 |
| qsub – submit pbs job                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | 9 |
| <b>SYNOPSIS</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | 9 |
| qsub [-a date_time] [-A account_string] [-c interval] [-C directive_prefix] [-e path] [-h] [-I] [-j<br>join] [-k keep] [-l resource_list] [-m mail_options] [-M user_list] [-N name] [-o path] [-p priori-<br>ty] [-q destination] [-r c] [-S path_list] [-u user_list] [-v variable_list] [-V] [-W additional_at-<br>tributes] [-z] [script]                                                                                                                                                                                                                                                                                                                                                                     | 9 |
| <b>DESCRIPTION</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | 9 |
| To create a job is to submit an executable script to a batch server. The batch server will be<br>the default server unless the <code>-q</code> option is specified. See discussion of <code>PBS_DEFAULT</code> under<br>Environment Variables below. Typically, the script is a shell script which will be executed by<br>a command shell such as <code>sh</code> or <code>csch</code> .                                                                                                                                                                                                                                                                                                                          | 9 |
| Options on the <b>qsub</b> command allow the specification of attributes which affect the behavior<br>of the job. The job is created by sending a <i>Queue Job</i> batch request to the batch server.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | 9 |
| The <code>qsub</code> command will pass certain environment variables in the <code>Variable_List</code> attribute of<br>the job. These variables will be available to the job. The value for the following variables<br>will be taken from the environment of the <code>qsub</code> command: <b>HOME</b> , <b>LANG</b> , <b>LOGNAME</b> ,<br><b>PATH</b> , <b>MAIL</b> , <b>SHELL</b> , and <b>TZ</b> . These values will be assigned to a new name which is the<br>current name prefixed with the string "PBS_O_". For example, the job will have access to an<br>environment variable named <b>PBS_O_HOME</b> which have the value of the variable <b>HOME</b> in<br>the <code>qsub</code> command environment. | 9 |
| In addition to the above, the following environment variables will be available to the batch<br>job. (The values of the following environment variables are established by <code>qsub</code> .)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | 9 |
| <b>PBS_O_HOST</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | 9 |
| the name of the host upon which the <code>qsub</code> command is running.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | 9 |
| <b>PBS_O_QUEUE</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | 9 |
| the name of the original queue to which the job was submitted. (It is established by the<br>server which creates the job, not <code>qsub</code> .)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | 9 |
| <b>PBS_O_SYSTEM</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | 9 |
| the operating system name given by <code>uname -s</code> on the host on which <code>qsub</code> is running.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | 9 |
| <b>PBS_O_WORKDIR</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | 9 |
| the absolute path of the current working directory of the <code>qsub</code> command.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | 9 |
| The following are established by the server executing the job, not the <code>qsub</code> command.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | 9 |
| <b>PBS_ENVIRONMENT</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | 9 |
| set to <code>PBS_BATCH</code> to indicate the job is a batch job, or to <code>PBS_INTERACTIVE</code> to indi-<br>cate the job is a PBS interactive job, see <code>-I</code> option.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | 9 |
| <b>PBS_JOBID</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | 9 |
| the job identifier assigned to the job by the batch system.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | 9 |
| <b>PBS_JOBNAME</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | 9 |
| the job name supplied by the user.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | 9 |
| <b>PBS_NODEFILE</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | 9 |
| the name of the file contain the list of nodes assigned to the job.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | 9 |
| <b>PBS_QUEUE</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | 9 |
| the name of the queue from which the job is executed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | 9 |
| <b>OPTIONS</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | 9 |

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |   |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| -a <i>date_time</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | 9 |
| Declares the time after which the job is eligible for execution.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | 9 |
| The <i>date_time</i> argument is in the form: <i>[[[CC]YY]MM]DD]hhmm[.SS]</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | 9 |
| Where <i>CC</i> is the first two digits of the year (the century), <i>YY</i> is the second two digits of the year, <i>MM</i> is the two digits for the month, <i>DD</i> is the day of the month, <i>hh</i> is the hour, <i>mm</i> is the minute, and the optional <i>SS</i> is the seconds.                                                                                                                                                                                                                                                                                  | 9 |
| If the month, <i>MM</i> , is not specified, it will default to the current month if the specified day <i>DD</i> , is in the future. Otherwise, the month will be set to next month. Likewise, if the day, <i>DD</i> , is not specified, it will default to today if the time <i>hhmm</i> is in the future. Otherwise, the day will be set to tomorrow. For example, if you submit a job at 11:15am with a time of <i>-a 1110</i> , the job will be eligible to run at 11:10am tomorrow. See the <i>date_time</i> operand for the <i>touch(1)</i> command defined by POSIX.2. | 9 |
| The <i>Execution_Time</i> job attribute will be set to the number of seconds since Epoch which is equivalent to the Universal time expressed by the local time in the <i>date_time</i> argument. If the <i>-a</i> option is not specified, the <i>Execution_Time</i> attribute is unset which represents a time zero or no delay.                                                                                                                                                                                                                                            | 9 |
| -A <i>account_string</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | 9 |
| Defines the account string associated with the job. The <i>account_string</i> is an undefined string of characters and is interpreted by the server which executes the job. See section 2.7.1 of the PBS ERS. The <i>Account_Name</i> attribute is set to the account string. If <i>account_string</i> is unset, it is not passed with the job to the job executor.                                                                                                                                                                                                          | 9 |
| -c <i>interval</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | 9 |
| Defines the interval at which the job will be checkpointed. If the job executes upon a host which does not support checkpoint, this option will be ignored.                                                                                                                                                                                                                                                                                                                                                                                                                  | 9 |
| The <i>interval</i> argument is specified as:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | 9 |
| n No checkpointing is to be performed. The job's <i>Checkpoint</i> attribute is set to the string "n".                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | 9 |
| s Checkpointing is to be performed only when the server executing the job is shut-down. The job's <i>Checkpoint</i> attribute is set to the string "s".                                                                                                                                                                                                                                                                                                                                                                                                                      | 9 |
| c Checkpointing is to be performed at the default minimum time for the server executing the job. The job's <i>Checkpoint</i> attribute is set to the string "c".                                                                                                                                                                                                                                                                                                                                                                                                             | 9 |
| c= <i>minutes</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | 9 |
| Checkpointing is to be performed at an interval of <i>minutes</i> , which is the integer number of minutes of CPU time used by the job. This value must be greater than zero. The <i>Checkpoint</i> attribute is set to the string specified by "c= <i>minutes</i> ".                                                                                                                                                                                                                                                                                                        | 9 |
| u Checkpointing is unspecified. Unless otherwise stated, "u" is treated the same as "s".                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | 9 |
| If <i>-c</i> is not specified, the <i>Checkpoint</i> attribute is set to the value "u".                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | 9 |
| -C <i>directive_prefix</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | 9 |
| Defines the prefix that declares a directive to the <i>qsub</i> command within the script file. See the paragraph on script directives in the Extended Description section.                                                                                                                                                                                                                                                                                                                                                                                                  | 9 |
| If the <i>-C</i> option is presented with a <i>directive_prefix</i> argument that is the null string, <i>qsub</i> will not scan the script file for directives. The directive prefix is not a job attribute. It is used solely within the <i>qsub</i> command.                                                                                                                                                                                                                                                                                                               | 9 |
| -e <i>path</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | 9 |
| Defines the path to be used for the standard error stream of the batch job. The <i>path</i> argument is of the form:                                                                                                                                                                                                                                                                                                                                                                                                                                                         | 9 |
| <i>[hostname:]path_name</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | 9 |
| where <i>hostname</i> is the name of a host to which the file will be returned and <i>path_name</i> is the path name on that host in the syntax recognized by POSIX. The argument will be interpreted as follows:                                                                                                                                                                                                                                                                                                                                                            | 9 |

|         |                                                                                                                                                                                                                                                                                                                                                                                                                                               |   |
|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
|         | path_name                                                                                                                                                                                                                                                                                                                                                                                                                                     | 9 |
|         | Where path_name is not an absolute path name, then the qsub command will expand the path name relative to the current working directory of the command. The command will supply the name of the host upon which it is executing for the <i>hostname</i> component.                                                                                                                                                                            | 9 |
|         | hostname:path_name                                                                                                                                                                                                                                                                                                                                                                                                                            | 9 |
|         | Where path_name is not an absolute path name, then the qsub command will not expand the path name relative to the current working directory of the command. On delivery of the standard error, the path name will be expanded relative to the user's home directory on the <i>hostname</i> system.                                                                                                                                            | 9 |
|         | path_name                                                                                                                                                                                                                                                                                                                                                                                                                                     | 9 |
|         | Where path_name specifies an absolute path name, then the qsub will supply the name of the host on which it is executing for the <i>hostname</i> .                                                                                                                                                                                                                                                                                            | 9 |
|         | hostname:path_name                                                                                                                                                                                                                                                                                                                                                                                                                            | 9 |
|         | Where path_name specifies an absolute path name, the path will be used as specified.                                                                                                                                                                                                                                                                                                                                                          | 9 |
|         | If the <i>-e</i> option is not specified, the default file name for the standard error stream will be used. The default name has the following form:                                                                                                                                                                                                                                                                                          | 9 |
|         | job_name.e <sub>sequence_number</sub>                                                                                                                                                                                                                                                                                                                                                                                                         | 9 |
|         | where <i>job_name</i> is the name of the job, see <i>-N</i> option, and <i>sequence_number</i> is the job number assigned when the job is submitted. This option sets the job attribute <i>Error_Path</i> .                                                                                                                                                                                                                                   | 9 |
| -h      | Specifies that a user hold be applied to the job at submission time. The <i>Hold_Types</i> attribute will be set to <i>USER</i> , "u". If <i>-h</i> is not specified, then <i>Hold_Types</i> is set to <i>NONE</i> , "n".                                                                                                                                                                                                                     | 9 |
| -I      | Declares that the job is to be run "interactively". The job will be queued and scheduled as any PBS batch job, but when executed, the standard input, output, and error streams of the job are connected through qsub to the terminal session in which qsub is running. See the "Extended Description" paragraph for addition information of interactive jobs.                                                                                | 9 |
| -j join | Declares if the standard error stream of the job will be merged with the standard output stream of the job.                                                                                                                                                                                                                                                                                                                                   | 9 |
|         | An option argument value of <i>oe</i> directs that the two streams will be merged, intermixed, as standard output. The <i>Join_Path</i> job attribute is set to "oe". An option argument value of <i>eo</i> directs that the two streams will be merged, intermixed, as standard error. The <i>Join_Path</i> job attribute is set to "eo".                                                                                                    | 9 |
|         | If the <i>join</i> argument is <i>n</i> or the option is not specified, the two streams will be two separate files. The <i>Join_Path</i> job attribute is set to "n".                                                                                                                                                                                                                                                                         | 9 |
| -k keep | Defines which (if either) of standard output or standard error will be retained on the execution host. If set for a stream, this option overrides the path name for that stream. If not set, neither stream is retained on the execution host.                                                                                                                                                                                                | 9 |
|         | The argument is either the single letter "e" or "o", or the letters "e" and "o" combined in either order. Or the argument is the letter <i>n</i> . Repetition of characters is permitted, but "n" may not appear in the same option argument with the other two characters. The attribute <i>Keep_Files</i> is set to the argument.                                                                                                           | 9 |
| e       | The standard error stream is to retained on the execution host. The stream will be placed in the home directory of the user under whose user id the job executed. The file name will be the default file name given by: <i>job_name.e<sub>sequence</sub></i> where <i>job_name</i> is the name specified for the job, and <i>sequence</i> is the sequence number component of the job identifier. The attribute is set to <i>KEEP_ERROR</i> . | 9 |

- o The standard output stream is retained on the execution host. The stream will be placed in the home directory of the user under whose user id the job executed. The file name will be the default file name given by: `job_name.osequence` where `job_name` is the name specified for the job, and `sequence` is the sequence number component of the job identifier. The attribute is set to `KEEP_OUTPUT`.
  - eo Both the standard output and standard error streams will be retained. The attribute is set to `"KEEP_OUTPUT | KEEP_ERROR"`.
  - oe Both the standard output and standard error streams will be retained. The attribute is set to `"KEEP_OUTPUT | KEEP_ERROR"`.
  - n Neither stream is retained.
- l resource\_list**  
 Defines the resources that are required by the job and establishes a limit to the amount of resource that can be consumed. If not set for a generally available resource, such as CPU time, the limit is infinite. The *resource\_list* argument is of the form:  
*resource\_name*=[*value*][,*resource\_name*=[*value*],...]  
 For each resource listed in the *resource\_list*, one entry will be added to the `Resource_List` attribute of the job. The entry contains the resource name and its requested value. No white space is allowed in the value. Other than syntax, `qsub` performs no resource or value checking. The checking is performed by the execution server.
- m mail\_options**  
 Defines the set of conditions under which the execution server will send a mail message about the job. The *mail\_options* argument is a string which consists of either the single character "n", or one or more of the characters "a", "b", and "e". Repeated letters are accepted, but n cannot be mixed with the other characters.  
 If the character "n" is specified, no mail will be sent. The `Mail_Points` attribute is set to `NONE`, "n".  
 For the letters "a", "b", and "e":  
 a mail is sent when the job is aborted by the batch system. The `Mail_Points` attribute is set to `ABORT`, "a".  
 b mail is sent when the job begins execution. The `Mail_Points` attribute is set to `BEGINNING`, "b".  
 e mail is sent when the job terminates. The `Mail_Points` attribute is set to `EXIT`, "e".  
 If the `-m` option is not specified, mail will be sent if the job is aborted. The `Mail_Points` attribute is set to `ABORT`, "a".
- M user\_list**  
 Declares the list of users to whom mail is sent by the execution server when it sends mail about the job.  
 The *user\_list* argument is of the form:  
*user*[*@host*][,*user*[*@host*],...]  
 If unset, the list defaults to the submitting user at the `qsub` host, i.e. the job owner.  
 The `Mail_Users` attribute is set to the argument.
- N name**  
 Declares a name for the job. The name specified may be up to and including 15 characters in length. It must consist of printable, non white space characters with the first character alphabetic. [The POSIX 1003.2d Standard calls for only alphanumeric characters, but then calls for the use of the script file base name as the job name if a name is not specified. The file name may contain other than alphanumeric

- ic characters. Therefore I “interpret” the standard as allowing printable characters.] 9  
 Names taken from the script name may have a non-alphabetic character first. If the 9  
 script basename is greater than 15 characters, it will be truncated to 15. 9
- If the *-N* option is not specified, the job name will be the base name of the job script 9  
 file specified on the command line. If no script file name was specified and the script 9  
 was read from the standard input, then the job name will be set to *STDIN*. 9
- The Job\_Name attribute is set to the name. 9
- o path** Defines the path to be used for the standard output stream of the batch job. The 9  
*path* argument is of the form: 9  
*[hostname:]path\_name* 9  
 where *hostname* is the name of a host to which the file will be returned and 9  
*path\_name* is the path name on that host in the syntax recognized by POSIX. The 9  
 argument will be interpreted as follows: 9
- path\_name* 9  
 Where *path\_name* is not an absolute path name, then the qsub command will 9  
 expand the path name relative to the current working directory of the com- 9  
 mand. The command will supply the name of the host upon which it is execut- 9  
 ing for the *hostname* component. 9
- hostname:path\_name* 9  
 Where *path\_name* is not an absolute path name, then the qsub command will 9  
 not expand the path name relative to the current working directory of the com- 9  
 mand. On delivery of the standard output, the path name will be expanded 9  
 relative to the user’s home directory on the *hostname* system. 9
- path\_name* 9  
 Where *path\_name* specifies an absolute path name, then the qsub will supply 9  
 the name of the host on which it is executing for the *hostname*. 9
- hostname:path\_name* 9  
 Where *path\_name* specifies an absolute path name, the path will be used as 9  
 specified. 9
- If the *-o* option is not specified, the default file name for the standard output stream 9  
 will be used. The default name has the following form: 9  
*job\_name.osequence\_number* 9  
 where *job\_name* is the name of the job, see *-N* option, and *sequence\_number* is 9  
 the job number assigned when the job is submitted. This option sets the job at- 9  
 tribute Output\_Path. 9
- p priority** 9  
 Defines the priority of the job. The *priority* argument must be a integer between 9  
 -1024 and +1023 inclusive. The default is no priority which is equivalent to a priori- 9  
 ty of zero. The Priority job attribute is set to this signed integer value. 9
- q destination** 9  
 Defines the destination of the job. The *destination* names a queue, a server, or a 9  
 queue at a server. 9  
 The qsub command will submit the script to the server defined by the *destination* 9  
 argument. The server named by the destination is the one to which qsub sends the 9  
*Queue Job* batch request. If the destination is a *routing queue*, the job may be rout- 9  
 ed by the server to a new destination. 9
- If the *-q* option is not specified, the qsub command will submit the script to the de- 9  
 fault server. See PBS\_DEFAULT under the Environment Variables section on this 9  
 man page and the PBS ERS section 2.7.4, "Default Server". 9
- If the *-q* option is specified, it is in one of the following three forms: 9  
*queue* 9



|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |   |
|--------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
|                          | <i>@server</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | 9 |
|                          | <i>queue@server</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | 9 |
|                          | If the <i>destination</i> argument names a queue and does not name a server, the job will be submitted to the named queue at the default server.                                                                                                                                                                                                                                                                                                                                                                                                                                        | 9 |
|                          | If the <i>destination</i> argument names a server and does not name a queue, the job will be submitted to the default queue at the named server.                                                                                                                                                                                                                                                                                                                                                                                                                                        | 9 |
|                          | If the <i>destination</i> argument names both a queue and a server, the job will be submitted to the named queue at the named server.                                                                                                                                                                                                                                                                                                                                                                                                                                                   | 9 |
| -r y n                   | Declares whether the job is rerunable. See the <b>qrerun</b> command. The option argument is a single character, either <i>y</i> or <i>n</i> . Also see <i>rerunable</i> in the glossary.                                                                                                                                                                                                                                                                                                                                                                                               | 9 |
|                          | If the argument is "y", the job is rerunable. The <i>Rerunable</i> attribute is set to the character 'y'. If the argument is "n", the job is not rerunable. The default value is 'y', rerunable.                                                                                                                                                                                                                                                                                                                                                                                        | 9 |
| -S path_list             | Declares the shell that interprets the job script.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | 9 |
|                          | The option argument <i>path_list</i> is in the form:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | 9 |
|                          | <i>path[@host][/,path[@host],...]</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | 9 |
|                          | Only one path may be specified for any host named. Only one path may be specified without the corresponding host name. The path selected will be the one with the host name that matched the name of the execution host. If no matching host is found, then the path specified without a host will be selected, if present.                                                                                                                                                                                                                                                             | 9 |
|                          | If the <i>-S</i> option is not specified, the option argument is the null string, or no entry from the <i>path_list</i> is selected, the execution will use the user's login shell on the execution host. The <i>Shell_Path_List</i> attribute is set to the <i>path_list</i> argument if present, otherwise it is set to the null string.                                                                                                                                                                                                                                              | 9 |
| -u user_list             | Defines the user name under which the job is to run on the execution system. If unset, the user list defaults to the user who is running qsub.                                                                                                                                                                                                                                                                                                                                                                                                                                          | 9 |
|                          | The <i>user_list</i> argument is of the form:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | 9 |
|                          | <i>user[@host][/,user[@host],...]</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | 9 |
|                          | Only one user name may be given per specified host. Only one of the <i>user</i> specifications may be supplied without the corresponding <i>host</i> specification. That user name will be used for execution on any host not named in the argument list. A named host refers to the host on which the job is queued for execution, not the actual execution host. Authorization must exist for the job owner to run as the specified user. See the <i>ruserok(3)</i> function for authorization requirements. The <i>User_List</i> attribute is set to the value of <i>user_list</i> . | 9 |
| -v variable_list         | Expands the list of environment variables that are exported to the job.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | 9 |
|                          | In addition to the variables described in the "Description" section above, <i>variable_list</i> names environment variables from the qsub command environment which are made available to the job when it executes. The <i>variable_list</i> is a comma separated list of strings of the form <i>variable</i> or <i>variable=value</i> . These variables and their values are passed to the job. The <i>Variable_List</i> attribute is appended with the variables in <i>user_list</i> and their values.                                                                                | 9 |
| -V                       | Declares that all environment variables in the qsub command's environment are to be exported to the batch job. The <i>Variable_List</i> attribute is appended with the variables in the qsub command's environment and their values.                                                                                                                                                                                                                                                                                                                                                    | 9 |
| -W additional_attributes | The <i>-W</i> option allows for the specification of additional job attributes. The general                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | 9 |

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |   |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| syntax of the <code>-W</code> is in the form:                                                                                                                                                                                                                                                                                                                                                                                                                            | 9 |
| <code>-W attr_name=attr_value[,attr_name=attr_value...]</code>                                                                                                                                                                                                                                                                                                                                                                                                           | 9 |
| Note if white space occurs anywhere within the option argument string or the equal sign, "=", occurs within an <i>attribute_value</i> string, then the string must be enclosed with either single or double quote marks.                                                                                                                                                                                                                                                 | 9 |
| PBS currently supports the following attributes within the <code>-W</code> option.                                                                                                                                                                                                                                                                                                                                                                                       | 9 |
| <i>depend=dependency_list</i>                                                                                                                                                                                                                                                                                                                                                                                                                                            | 9 |
| Defines the dependency between this and other jobs. The <i>dependency_list</i> is in the form:                                                                                                                                                                                                                                                                                                                                                                           | 9 |
| <i>type[:argument[:argument...]][,type:argument...]</i> .                                                                                                                                                                                                                                                                                                                                                                                                                | 9 |
| The <i>argument</i> is either a numeric count or a PBS job id according to <i>type</i> . If argument is a count, it must be greater than 0. If it is a job id and not fully specified in the form <i>seq_number.server.name</i> , it will be expanded according to the default server rules which apply to job ids on most commands. If <i>argument</i> is null (the preceding colon need not be specified), the dependency of the coresponding type is cleared (unset). | 9 |
| <i>synccount:count</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                   | 9 |
| This job is the first in a set of jobs to be executed at the same time. <i>Count</i> is the number of additional jobs in the set.                                                                                                                                                                                                                                                                                                                                        | 9 |
| <i>syncwith:jobid</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                    | 9 |
| This job is an additional member of a set of jobs to be executed at the same time. In the above and following dependency types, <i>jobid</i> is the job identifier of the first job in the set.                                                                                                                                                                                                                                                                          | 9 |
| <i>after:jobid[:jobid...]</i>                                                                                                                                                                                                                                                                                                                                                                                                                                            | 9 |
| This job may be scheduled for execution at any point after jobs <i>jobid</i> have started execution.                                                                                                                                                                                                                                                                                                                                                                     | 9 |
| <i>afterok:jobid[:jobid...]</i>                                                                                                                                                                                                                                                                                                                                                                                                                                          | 9 |
| This job may be scheduled for execution only after jobs <i>jobid</i> have terminated with no errors. See the csh warning under "Extended Description".                                                                                                                                                                                                                                                                                                                   | 9 |
| <i>afternotok:jobid[:jobid...]</i>                                                                                                                                                                                                                                                                                                                                                                                                                                       | 9 |
| This job may be scheduled for execution only after jobs <i>jobid</i> have terminated with errors. See the csh warning under "Extended Description".                                                                                                                                                                                                                                                                                                                      | 9 |
| <i>afterany:jobid[:jobid...]</i>                                                                                                                                                                                                                                                                                                                                                                                                                                         | 9 |
| This job may be scheduled for execution after jobs <i>jobid</i> have terminated, with or without errors.                                                                                                                                                                                                                                                                                                                                                                 | 9 |
| <i>on:count</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                          | 9 |
| This job may be scheduled for execution after <i>count</i> dependencies on other jobs have been satisfied. This form is used in conjunction with one of the <i>before</i> forms, see below.                                                                                                                                                                                                                                                                              | 9 |
| <i>before:jobid[:jobid...]</i>                                                                                                                                                                                                                                                                                                                                                                                                                                           | 9 |
| When this job has begun execution, then jobs <i>jobid...</i> may begin.                                                                                                                                                                                                                                                                                                                                                                                                  | 9 |
| <i>beforeok:jobid[:jobid...]</i>                                                                                                                                                                                                                                                                                                                                                                                                                                         | 9 |
| If this job terminates execution without errors, then jobs <i>jobid...</i> may begin. See the csh warning under "Extended Description".                                                                                                                                                                                                                                                                                                                                  | 9 |
| <i>beforenotok:jobid[:jobid...]</i>                                                                                                                                                                                                                                                                                                                                                                                                                                      | 9 |
| If this job terminates execution with errors, then jobs <i>jobid...</i> may begin. See the csh warning under "Extended Description".                                                                                                                                                                                                                                                                                                                                     | 9 |
| <i>beforeany:jobid[:jobid...]</i>                                                                                                                                                                                                                                                                                                                                                                                                                                        | 9 |
| When this job terminates execution, jobs <i>jobid...</i> may begin.                                                                                                                                                                                                                                                                                                                                                                                                      | 9 |

- If any of the *before* forms are used, the jobs referenced by *jobid* must have been submitted with a dependency type of *on*. 9
- The *depend* attribute is set to the value of the *dependency* option argument. 9
- If any of the *before* forms are used, the jobs referenced by *jobid* must have the same owner as the job being submitted. Otherwise, the dependency is ignored. 9
- Error processing of the existence, state, or condition of the job on which the newly submitted job is a deferred service, i.e. the check is performed after the job is queued. If an error is detected, the new job will be deleted by the server. Mail will be sent to the job submitter stating the error. 9
- Dependency examples: 9
- ```
qsub -W depend=afterok:123.big.iron.com /tmp/script 9
```
- ```
qsub -W depend=before:234.hunk1.com:235.hunk1.com /tmp/script 9
```
- ```
group_list=g_list 9
```
- Defines the group name under which the job is to run on the execution system. The *g\_list* argument is of the form: 9
- ```
group[@host][/,group[@host],...] 9
```
- Only one group name may be given per specified host. Only one of the *group* specifications may be supplied without the corresponding *host* specification. That group name will be used for execution on any host not named in the argument list. The *group\_list* attribute is set to the value of *g\_list*. If not set, the *group\_list* defaults to the primary group of the user under which the job will be run. 9
- ```
interactive=true 9
```
- If the *interactive* attribute is specified, the job is an interactive job. The *-I* option is an alternative method of specifying this attribute. 9
- ```
stagein=file_list 9
```
- ```
stageout=file_list 9
```
- Specifies the *stagein* or *stageout* attribute, listing which files are staged (copied) in before job start or staged out after the job completes execution. On completion of the job, all staged-in and staged-out files are removed from the execution system. The *file\_list* is in the form 9
- ```
local_file@hostname:remote_file[...] 9
```
- regardless of the direction of the copy. The name *local\_file* is the name of the file on the system where the job executed. It may be an absolute path or relative to the home directory of the user. The name *remote\_file* is the destination name on the host specified by *hostname*. The name may be absolute or relative to the user's home directory on the destination host. The use of wildcards in the file name is not recommended. Since *rcp* (or *scp*) is run via *rsh*, it will pick up matching names from the remote system. However, *pbs\_mom* will not expand the wildcards and will fail to delete the staged files on job termination. The file names map to a remote copy program (*rcp*) call on the execution system in the following manner: 9
- ```
For stagein: rcp hostname:remote_file local_file 9
```
- ```
For stageout: rcp local_file hostname:remote_file 9
```
- z* Directs that the *qsub* command is not to write the job identifier assigned to the job to the command's standard output. 9

## OPERANDS 9

The *qsub* command accepts a *script* operand that is the path to the script of the job. If the path is relative, it will be expanded relative to the working directory of the *qsub* command. 9

If the *script* operand is not provided or the operand is the single character "-", the *qsub* command reads the script from standard input. 9

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |   |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| <b>STANDARD INPUT</b>                                                                                                                                                                                                                                                                                                                                                                                                                                           | 9 |
| The qsub command reads the script for the job from standard input if the <i>script</i> operand is missing or is the single character "-".                                                                                                                                                                                                                                                                                                                       | 9 |
| <b>INPUT FILES</b>                                                                                                                                                                                                                                                                                                                                                                                                                                              | 9 |
| The <i>script</i> file is read by the qsub command. Qsub acts upon any directives found in the script.                                                                                                                                                                                                                                                                                                                                                          | 9 |
| When the job is created, a copy of the script file is made and that copy cannot be modified.                                                                                                                                                                                                                                                                                                                                                                    | 9 |
| <b>STANDARD OUTPUT</b>                                                                                                                                                                                                                                                                                                                                                                                                                                          | 9 |
| Unless the <i>-z</i> option is set, the job identifier assigned to the job will be written to standard output if the job is successfully created.                                                                                                                                                                                                                                                                                                               | 9 |
| <b>STANDARD ERROR</b>                                                                                                                                                                                                                                                                                                                                                                                                                                           | 9 |
| The qsub command will write a diagnostic message to standard error for each error occurrence.                                                                                                                                                                                                                                                                                                                                                                   | 9 |
| <b>ENVIRONMENT VARIABLES</b>                                                                                                                                                                                                                                                                                                                                                                                                                                    | 9 |
| The values of some or all of the variables in the qsub command's environment are exported with the job, see the <i>-v</i> and <i>-V</i> options.                                                                                                                                                                                                                                                                                                                | 9 |
| The environment variable <b>PBS_DEFAULT</b> defines the name of the default server. Typically, it corresponds to the system name of the host on which the server is running. If <b>PBS_DEFAULT</b> is not set, the default is defined by an administrator established file.                                                                                                                                                                                     | 9 |
| The environment variable <b>PBS_DPREFIX</b> determines the prefix string which identifies directives in the script.                                                                                                                                                                                                                                                                                                                                             | 9 |
| <b>EXTENDED DESCRIPTION</b>                                                                                                                                                                                                                                                                                                                                                                                                                                     | 9 |
| Script Processing:                                                                                                                                                                                                                                                                                                                                                                                                                                              | 9 |
| A job script may consist of PBS directives, comments and executable statements. A PBS directive provides a way of specifying job attributes in addition to the command line options. For example:                                                                                                                                                                                                                                                               | 9 |
| <pre> : #PBS -N Job_name #PBS -l walltime=10:30,mem=320kb #PBS -m be # step1 arg1 arg2 step2 arg3 arg4 </pre>                                                                                                                                                                                                                                                                                                                                                   | 9 |
| The qsub command scans the lines of the script file for directives. An initial line in the script that begins with the characters "#!" or the character ":" will be ignored and scanning will start with the next line. Scanning will continue until the first executable line, that is a line that is not blank, not a directive line, nor a line whose first non white space character is "#". If directives occur on subsequent lines, they will be ignored. | 9 |
| A line in the script file will be processed as a directive to qsub if and only if the string of characters starting with the first non white space character on the line and of the same length as the directive prefix matches the directive prefix.                                                                                                                                                                                                           | 9 |
| The remainder of the directive line consists of the options to qsub in the same syntax as they appear on the command line. The option character is to be preceded with the "-" character.                                                                                                                                                                                                                                                                       | 9 |
| If an option is present in both a directive and on the command line, that option and its argument, if any, will be ignored in the directive. The command line takes precedence.                                                                                                                                                                                                                                                                                 | 9 |

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |   |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| If an option is present in a directive and not on the command line, that option and its argument, if any, will be processed as if it had occurred on the command line.                                                                                                                                                                                                                                                                                                                                                  | 9 |
| The directive prefix string will be determined in order of preference from:                                                                                                                                                                                                                                                                                                                                                                                                                                             | 9 |
| The value of the <code>-C</code> option argument if the option is specified on the command line.                                                                                                                                                                                                                                                                                                                                                                                                                        | 9 |
| The value of the environment variable <b>PBS_DPREFIX</b> if it is defined.                                                                                                                                                                                                                                                                                                                                                                                                                                              | 9 |
| The four character string <code>#PBS</code> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | 9 |
| If the <code>-C</code> option is found in a directive in the script file, it will be ignored.                                                                                                                                                                                                                                                                                                                                                                                                                           | 9 |
| User Authorization:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | 9 |
| When the user submits a job from a system other than the one on which the PBS Server is running, the name under which the job is to be executed is selected according to the rules listed under the <code>-u</code> option. The user submitting the job must be authorized to run the job under the execution user name. This authorization is provided if                                                                                                                                                              | 9 |
| (1) The host on which <code>qsub</code> is run is trusted by the execution host (see <code>/etc/hosts.equiv</code> ),                                                                                                                                                                                                                                                                                                                                                                                                   | 9 |
| (2) The execution user has an <code>.rhosts</code> file naming the submitting user on the submitting host.                                                                                                                                                                                                                                                                                                                                                                                                              | 9 |
| C-Shell <code>.logout</code> File:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | 9 |
| The following warning applies for users of the c-shell, <code>csh</code> . If the job is executed under the <code>csh</code> and a <code>.logout</code> file exists in the home directory in which the job executes, the exit status of the job is that of the <code>.logout</code> script, not the job script. This may impact any inter-job dependencies. To preserve the job exit status, either remove the <code>.logout</code> file or place the following line as the first line in the <code>.logout</code> file | 9 |
| <i>set EXITVAL = \$status</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | 9 |
| and the following line as the last executable line in <code>.logout</code>                                                                                                                                                                                                                                                                                                                                                                                                                                              | 9 |
| <i>exit \$EXITVAL</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | 9 |
| Interactive Jobs:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | 9 |
| If the <code>-I</code> option is specified on the command line or in a script directive, or if the "interactive" job attribute declared true via the <code>-W</code> option, <code>-W interactive=true</code> , either on the command line or in a script directive, the job is an interactive job. The script will be processed for directives, but will not be included with the job. When the job begins execution, all input to the job is from the terminal session in which <code>qsub</code> is running.         | 9 |
| When an interactive job is submitted, the <code>qsub</code> command will not terminate when the job is submitted. <code>Qsub</code> will remain running until the job terminates, is aborted, or the user interrupts <code>qsub</code> with an <code>SIGINT</code> (the control-C key). If <code>qsub</code> is interrupted prior to job start, it will query if the user wishes to exit. If the user response "yes", <code>qsub</code> exits and the job is aborted.                                                   | 9 |
| Once the interactive job has started execution, input to and output from the job pass through <code>qsub</code> . Keyboard generated interrupts are passed to the job. Lines entered that begin with the tilde ("~") character and contain special sequences are escaped by <code>qsub</code> . The recognized escape sequences are:                                                                                                                                                                                    | 9 |
| ~ <code>Qsub</code> terminates execution. The batch job is also terminated.                                                                                                                                                                                                                                                                                                                                                                                                                                             | 9 |
| ~susp Suspend the <code>qsub</code> program if running under the C shell. "susp" is the suspend character, usually <code>CNTL-Z</code> .                                                                                                                                                                                                                                                                                                                                                                                | 9 |
| ~asusp Suspend the input half of <code>qsub</code> (terminal to job), but allow output to continue to be displayed. Only works under the C shell. "asusp" is the auxiliary suspend character, usually <code>CNTL-Y</code> .                                                                                                                                                                                                                                                                                             | 9 |

|                                                                                          |   |
|------------------------------------------------------------------------------------------|---|
| <b>EXIT STATUS</b>                                                                       | 9 |
| Upon successful processing, the qsub exit status will be a value of zero.                | 9 |
| If the qsub command fails, the command exits with a value greater than zero.             | 9 |
| <b>SEE ALSO</b>                                                                          | 9 |
| qalter(1B), qdel(1B), qhold(1B), qmove(1B), qmsg(1B), qrerun(1B), qrls(1B), qselect(1B), | 9 |
| qsig(1B), qstat(1B), pbs_connect(3B),, ruserok(3), pbs_job_attributes(7B), pbs_queue_at- | 9 |
| tributes(7B), pbs_resources_irix5(7B), pbs_resources_sp2(7B), pbs_resources_sunos4(7B),  | 9 |
| pbs_resources_unicos8(7B), pbs_server_attributes(7B), and pbs_server(8B)                 | 9 |

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |   |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| 5.2.13. Convert NQS Scripts                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | 9 |
| <b>NAME</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | 9 |
| nqs2pbs – convert NQS job scripts to PBS                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | 9 |
| <b>SYNOPSIS</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | 9 |
| nqs2pbs nqs_script [pbs_script]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | 9 |
| <b>DESCRIPTION</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | 9 |
| This utility converts a existing NQS job script to work with PBS and NQS. The existing script is copied and PBS directives, <i>#PBS</i> , are inserted prior to each NQS directive <i>#QSUB</i> or <i>#@\$</i> , in the original script.                                                                                                                                                                                                                                                                             | 9 |
| Certain NQS date specification and options are not supported by PBS. A warning message will be displayed indicating the problem and the line of the script on which it occurred.                                                                                                                                                                                                                                                                                                                                     | 9 |
| If any unrecognizable NQS directives are encountered, an error message is displayed. The new PBS script will be deleted if any errors occur.                                                                                                                                                                                                                                                                                                                                                                         | 9 |
| <b>OPERANDS</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | 9 |
| nqs_script                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | 9 |
| Specifies the file name of the NQS script to convert. This file is not changed.                                                                                                                                                                                                                                                                                                                                                                                                                                      | 9 |
| pbs_script                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | 9 |
| If specified, it is the name of the new PBS script. If not specified, the new file name is <i>nqs_script.new</i> .                                                                                                                                                                                                                                                                                                                                                                                                   | 9 |
| <b>NOTES</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | 9 |
| Converting NQS date specifications to the PBS form may result in a warning message and an incompleted converted date. PBS does not support date specifications of "today", "tomorrow", or the name of the days of the week such as "Monday". If any of these are encountered in a script, the PBS specification will contain only the time portion of the NQS specification, i.e. <i>#PBS -a hhmm[.ss]</i> . It is suggested that you specify the execution time on the qsub command line rather than in the script. | 9 |
| Note that PBS will interpret a time specification without a date in the following way:                                                                                                                                                                                                                                                                                                                                                                                                                               | 9 |
| - If the time specified has not yet been reached, the job will become eligible to run at that time today.                                                                                                                                                                                                                                                                                                                                                                                                            | 9 |
| - If the specified time has already passed when the job is submitted, the job will become eligible to run at that time tomorrow.                                                                                                                                                                                                                                                                                                                                                                                     | 9 |
| PBS does not support time zone identifiers. All times are taken as local time.                                                                                                                                                                                                                                                                                                                                                                                                                                       | 9 |
| <b>SEE ALSO</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | 9 |
| qsub(1B)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | 9 |

## 5.2.14. BASL Compiler 9

**NAME** 9

`basl2c` – converts a BASL (BAtch Scheduling Language) code into a C scheduler code. 9

**SYNOPSIS** 9

`basl2c [-d] [-l lexerDebugFile] [-p parserDebugFile] [-y symtabDebugFile] [-s semanticDebugFile] [-g codegenDebugFile] [-c cFile] baslFile` 9

**DESCRIPTION** 9

`basl2c` is the BASL to C compiler that produces an intermediate code that can be fed into a regular C compiler, and linked with the PBS libraries to produce the scheduler executable. `Basl2c` takes as input a *baslFile*, which is a program written in the BAtch Scheduling Language, containing the main scheduling code. `Basl2c` then converts the BASL constructs in the file into C statements, and it also attaches additional code to produce the PBS scheduler source code. By default, the resulting C code is written into the file *pbs\_sched.c*. 9

The full pathname to the resulting C file is what needs to be specified in the **SCHD\_CODE** variable in `local.mk` before compiling the BASL scheduler to produce the *pbs\_sched* executable. 9

**OPTIONS** 9

`-d` Prints additional debugging messages to the lexer (see `-l` option), parser (see `-p` option), symbol table (see `-y` option), semantic analyzer (see `-s` option), and code generator (see `-g` option). 9

`-l lexerDebugFile`  
*lexerDebugFile* is the name of a file to write into the debugging messages generated while scanning for tokens. 9

`-p parserDebugFile`  
*parserDebugFile* is the name of a file to write into the debugging messages generated while putting together tokens in a usable way. 9

`-y symtabDebugFile`  
*symtabDebugFile* is the name of a file to write into the debugging messages related to the symbol table. 9

`-s semanticDebugFile`  
*semanticDebugFile* is the name of a file to write into the debugging messages generated while checking to make sure variables and operators are used in a consistent way. 9

`-g codegenDebugFile`  
*codegenDebugFile* is the name of a file to write into the debugging messages generated while converting BASL statements to C statements. 9

`-c cFile`  
*cFile* is the name of a file where the generated C code is written into. 9

**MAIN STRUCTURE** 9

The basic structure of a scheduler code written in BASL is as follows: 9



```

zero or more FUNCTIONS definitions 9
zero or more global VARIABLE DECLARATIONS 9
zero or more assignment statements (to initialize global variables) 9
sched_main() 9
{ 9
    one or more VARIABLE DECLARATIONS 9
    zero or more STATEMENTS 9
} 9

```

For example, 9

```

% cat sched.basl 9
Int sum(Int a, Int b) 9
{ 9
    Int s; 9
    s = a + b; 9
    return(s); 9
} 9
Int glob; 9
sched_main() 9
{ 9
    Int c; 9
    a = 3; 9
    b = 4; 9
    c = sum(a, b); 9
    print(c); 9
    glob = 5; 9
    print(glob); 9
} 9

```

*sched\_main()* is the function that gets called at every scheduling iteration. 9

## **FUNCTIONS** 9

To define a function that can be called in subsequent functions, the syntax is: 9

```

ReturnType function-name ( DATATYPE1 IDENTIFIER1, 9
    DATATYPE2 IDENTIFIER2, ... ) 9
{ 9
    one or more VARIABLE DECLARATIONS 9
    zero or more STATEMENTS 9
} 9

```

For example, 9

```

Void printStuff(Dayofweek dow, DateTime t, String str, 9
                Size sz, CNode cn) 9
{ 9
    print(dow); 9
    print(t); 9
    print(str); 9
    print(sz); 9
    print(cn); 9
} 9

```

Valid function **ReturnType** are: Void, Int, Float, Dayofweek, DateTime, String, Size, Server, Que, Job, CNode, Set Server, Set Que, Set Job, Set CNode. 9

Valid data types ( **DATATYPE1, DATATYPE2, ...** ) for the parameter identifiers are: Int, Float, Dayofweek, DateTime, String, Size, Server, Que, Job, CNode, Set Server, Set Que, Set Job, Set CNode, Range Int, Range Float, Range Dayofweek, Range DateTime, Range Size, Fun Int, Fun Float, Fun Void, Fun Dayofweek, Fun DateTime, Fun String, Fun Size, Fun Server, Fun Que, Fun Job, Fun CNode, Fun Set Server, Fun Set Que, Fun Set Job, Fun Set CNode. These data types will be discussed in the next topic. 9

Functions are invoked by their name and their arguments as in: 9

```
printStuff( MON, (5|1|1997@14:32:00), "sched begins",
            30gb, node );
```

*basl2c* will actually add a "basl\_" prefix to the function name given by the scheduler writer to minimize chance of name collision, which can result when the resulting C code is linked with the PBS, BASL libraries. For example, if you look at the generated C code for *printStuff*, you would see, 9

```
basl_printStuff( MON, (5|1|1997@14:32:00),
                "sched begins", 30gb, node );
```

As in C, all function calls must have been previously defined. The BASL compiler will check to make sure that arguments in the function call match up exactly (in terms of types) with the parameters in the function definition. 9

Two kinds of functions exist in BASL: user-defined functions and predefined functions. User-defined functions are those that the scheduler writer provided a definition for, while predefined functions are those that can immediately be called without a need for defining it. For a list of predefined functions, see section on **PREDEFINED FUNCTIONS**. 9

**VARIABLE DECLARATIONS** 9

Like in C, all variables in a BASL code must be explicitly declared before use. Those variables declared outside of any function are referred to as global variables, while variables that are declared within a function body are called local variables. Global variables are usable anywhere within the BASL code, while local variables are readable only within the function from which they were declared. 9

The syntax of a variable declaration is: 9

```
DATATYPE IDENTIFIER ; 9
```

where **DATATYPE** can be: Int, Float, Dayofweek, DateTime, String, Size, Server, Que, Job, CNode, Set Server, Set Que, Set Job, Set CNode, Range Int, Range Float, Range Dayofweek, Range DateTime, Range Size. 9

**DATA TYPE** 9

**Void** used for functions that don't return a value. 9

**Int** signed, whole numbers given in base 10. 9

Sample constants: 9

5, +1, -3, SUCCESS (=1), FAIL (=0), TRUE (=1), FALSE (=0) 9

|                        |                                                                                                                                                                                                                                                                                                                                                                                     |                                 |
|------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>Float</b>           | real numbers which are represented as doubles in the translated C code.<br>Sample constants: 4.3, +1.2, -2.6                                                                                                                                                                                                                                                                        | 9<br>9                          |
| <b>Dayofweek</b>       | constant values: SUN, MON, TUE, WED, THU, FRI, SAT, internally represented as integer valued constants with SUN=0, MON=1, and so on.                                                                                                                                                                                                                                                | 9<br>9                          |
| <b>DateTime</b>        | specify in one of 3 formats:                                                                                                                                                                                                                                                                                                                                                        | 9                               |
|                        | [1] (m d y) where 1 <= m <= 12, 1 <= d <= 31, 0 <= y, ex. (4 4 1997);                                                                                                                                                                                                                                                                                                               | 9                               |
|                        | [2] (hh:mm:ss) where 0 <= hh <= 23, 0 <= mm <= 59, 0 <= ss <= 61, ex. (12:01:00);                                                                                                                                                                                                                                                                                                   | 9                               |
|                        | [3] (m d y@hh:mm:ss), ex. (4 4 1997@12:01:00)                                                                                                                                                                                                                                                                                                                                       | 9                               |
|                        | During dates/times comparison, "now" time is substituted if the time portion is not given (format [1]); the "now" date is substituted if the date portion is not given (format [2]). Also, the full year portion must be given (i.e. 1997 instead of 97) in dates to avoid ambiguity.                                                                                               | 9<br>9<br>9<br>9                |
| <b>String</b>          | A string is enclosed in quotes (") and it can contain anything except another quote, a newline, and left and right parentheses.<br>Sample constants: "a sample string", NULLSTR                                                                                                                                                                                                     | 9<br>9<br>9                     |
| <b>Size</b>            | format: <integer><suffix> where suffix is a multiplier of the form: <multiplier><unit>:                                                                                                                                                                                                                                                                                             | 9<br>9                          |
|                        | <i>multiplier</i> <i>unit (bytes or words)</i>                                                                                                                                                                                                                                                                                                                                      | 9                               |
|                        | =====                                                                                                                                                                                                                                                                                                                                                                               | 9                               |
|                        | <i>k,m,g,t,p,K,M,G,T,P</i> <i>b,B,w,W</i>                                                                                                                                                                                                                                                                                                                                           | 9                               |
|                        | where k=K=1024, m=M=1,048,576, g=G=1,073,741,824, t=T=1,099,511,627,776, p=P=1,125,899,906,842,624, b=B=1, and word size w=W is locally defined (i.e. 4 bytes in a 32-bit machine).                                                                                                                                                                                                 | 9<br>9                          |
|                        | When operating on 2 size operands that are of different suffixes, the suffix of the "lower" of the two will be the resultant suffix. For example,<br>10mb + 10gb = 10250mb                                                                                                                                                                                                          | 9<br>9<br>9                     |
|                        | Sample constants: -1b, 2w, 1kb, 2mw, +3gb, 4tw, 6Pb                                                                                                                                                                                                                                                                                                                                 | 9                               |
| <b>Range Int</b>       | format: (low Int value, high Int value)<br>where low Int value <= high Int value. Sample constant: (1,3)                                                                                                                                                                                                                                                                            | 9<br>9<br>9                     |
| <b>Range Float</b>     | format: (low Float value, high Float value)<br>where low value <= high value. Sample constant: (2.3, 4.6)                                                                                                                                                                                                                                                                           | 9<br>9<br>9                     |
| <b>Range Dayofweek</b> | format: (earlier day, later day)<br>where earlier day <= later day. Sample constant: (WED, FRI)                                                                                                                                                                                                                                                                                     | 9<br>9<br>9                     |
| <b>Range DateTime</b>  | format: (earlier date/time, later date/time)<br>where earlier date/time <= later date/time.<br>NOTE: if range contains only time portions, and earlier time "appears" to be > later time as in "((18:0:0), (6:0:0))", then during date/time comparisons, the "later" time will be adjusted by one day so that it will look like: "( (<now date>@18:0:0), (<tomorrow date>@6:0:0) )" | 9<br>9<br>9<br>9<br>9<br>9<br>9 |

|                                                                                                               |   |
|---------------------------------------------------------------------------------------------------------------|---|
| Sample constants:                                                                                             | 9 |
| ((4 4 1997), (4 10 1997)), ((12:01:00), (12:30:00)), ((4 4 1997@12:01:00),                                    | 9 |
| (4 10 1997@12:30:00))                                                                                         | 9 |
| <b>Range Size</b>                                                                                             | 9 |
| format: (low size, high size)                                                                                 | 9 |
| where low size <= high size. Sample constants: (23gb, 50gb)                                                   | 9 |
| <b>Server</b> Maps directly to the PBS server object. A <b>Server</b> manages one or more <b>Que</b> objects. | 9 |
| Sample constant: NOSERVER                                                                                     | 9 |
| <b>CNode</b> for computational node consisting of a single OS image, a shared memory, and a set               | 9 |
| of cpus. CNode runs 1 PBS MOM.                                                                                | 9 |
| Sample constant: NOCNODE                                                                                      | 9 |
| <b>Que</b> Maps directly to the PBS queue object. A <b>Que</b> object spools one or more <b>Job</b> objects.  | 9 |
| Sample constant: NOQUE                                                                                        | 9 |
| <b>Job</b> Maps directly to the PBS job object. A <b>Job</b> object carries some attributes and re-           | 9 |
| source requirements.                                                                                          | 9 |
| Sample constant: NOJOB                                                                                        | 9 |
| <b>Set Server</b>                                                                                             | 9 |
| list of Server objects.                                                                                       | 9 |
| Sample constant: EMPTYSETSERVER                                                                               | 9 |
| <b>Set CNode</b>                                                                                              | 9 |
| list of CNode objects.                                                                                        | 9 |
| Sample constant: EMPTYSETCNODE                                                                                | 9 |
| <b>Set Que</b> list of Que objects.                                                                           | 9 |
| Sample constant: EMPTYSETQUE                                                                                  | 9 |
| <b>Set Job</b> list of Job objects.                                                                           | 9 |
| Sample constant: EMPTYSETJOB                                                                                  | 9 |
| <br><b>BASL-DEFINED CONSTANTS</b>                                                                             | 9 |
| These are constants that cannot be used for naming an identifier (see next topic).                            | 9 |
| These are always in uppercase.                                                                                | 9 |
| <br><i>DATA TYPE</i> <i>BASL-DEFINED CONSTANT</i>                                                             | 9 |
| =====                                                                                                         | 9 |
| <i>Dayofweek</i> <i>SUN, MON, TUE, WED, THU, FRI, SAT</i>                                                     | 9 |
| <br><i>Int</i> <i>SUCCESS, FAIL, FALSE, TRUE, SYNCRUN, ASYNCRUN,</i>                                          | 9 |
| <i>DELETE, RERUN, HOLD, RELEASE, SIGNAL,</i>                                                                  | 9 |
| <i>MODIFYATTR, MODIFYRES, SERVER_ACTIVE,</i>                                                                  | 9 |
| <i>SERVER_IDLE, SERVER_SCHED, SERVER_TERM,</i>                                                                | 9 |
| <i>SERVER_TERMDELAY, QTYPE_E, QTYPE_R,</i>                                                                    | 9 |
| <i>SCHED_DISABLED, SCHED_ENABLED, TRANSIT,</i>                                                                | 9 |
| <i>QUEUED, HELD, WAITING, RUNNING, EXITING,</i>                                                               | 9 |
| <i>CNODE_OFFLINE, CNODE_DOWN, CNODE_FREE,</i>                                                                 | 9 |
| <i>CNODE_RESERVE, CNODE_INUSE_EXCLUSIVE,</i>                                                                  | 9 |
| <i>CNODE_INUSE_SHARED, CNODE_TIMESHARED,</i>                                                                  | 9 |
| <i>CNODE_CLUSTER, CNODE_UNKNOWN, OP_EQ, OP_NEQ,</i>                                                           | 9 |
| <i>OP_LE, OP_LT, OP_GE, OP_GT, OP_MAX, OP_MIN,</i>                                                            | 9 |
| <i>ASC, DESC</i>                                                                                              | 9 |
| <br><i>Server</i> <i>NOSERVER</i>                                                                             | 9 |

|                   |                       |   |
|-------------------|-----------------------|---|
| <i>Set Server</i> | <i>EMPTYSETSERVER</i> | 9 |
| <i>CNode</i>      | <i>NOCNODE</i>        | 9 |
| <i>Set CNode</i>  | <i>EMPTYSETCNODE</i>  | 9 |
| <i>Que</i>        | <i>NOQUE</i>          | 9 |
| <i>Set Que</i>    | <i>EMPTYSETQUE</i>    | 9 |
| <i>Job</i>        | <i>NOJOB</i>          | 9 |
| <i>Set Job</i>    | <i>EMPTYSETJOB</i>    | 9 |
| <i>String</i>     | <i>NULLSTR</i>        | 9 |

**IDENTIFIER** 9  
 Identifiers (used for variable names and function names) are in alphanumeric format, with the special underscore (`_`) character allowed. Currently, BASL can only handle identifiers with length of up to 80 chars. Also, you cannot use the BASL-defined constant names for naming an identifier. 9

**STATEMENTS** 9  
 In BASL(2), you can have a single statement terminated by a semi-colon, or a group of statements (called compound statement or block) delimited by '{' and '}'. The different kinds of statements that can appear in a BASL code are: 9

1. expression statement 9  
 Expression statements are anything of the form: 9

**expr ;** 9

where **expr** can be: 9

a) Arithmetic expressions 9

- lexpr + rexr (add)** 9
- lexpr - rexr (subtract)** 9
- lexpr \* rexr (multiply)** 9
- lexpr / rexr (divide)** 9
- lexpr % rexr (modulus or remainder)** 9

NOTE: Adding, subtracting, multiplying, dividing, and remaindering will only be allowed for proper types and if the left and right expressions are of consistent types. The table below illustrates what types are consistent among the various operators: 9

For +: 9

|                     |                                 |   |
|---------------------|---------------------------------|---|
| <i>lexpr</i>        | <i>rexpr</i>                    | 9 |
| =====               | =====                           | 9 |
| <i>Int or Float</i> | <i>Int or Float</i>             | 9 |
| <i>Size</i>         | <i>Size</i>                     | 9 |
| <i>String</i>       | <i>String, Int, Float, Size</i> | 9 |

For -, \*, /: 9

|              |              |   |
|--------------|--------------|---|
| <i>lexpr</i> | <i>rexpr</i> | 9 |
|--------------|--------------|---|

|                                                         |   |
|---------------------------------------------------------|---|
| =====                                                   | 9 |
| <i>Int or Float Int or Float</i>                        | 9 |
| <i>Size Size</i>                                        | 9 |
| For %:                                                  | 9 |
| <i>lexpr rexpr</i>                                      | 9 |
| =====                                                   | 9 |
| <i>Int or Float Int or Float</i>                        | 9 |
| Here are some sample arithmetic expressions statements: | 9 |
| <i>Int i1;</i>                                          | 9 |
| <i>Int i2;</i>                                          | 9 |
| <i>Float f1;</i>                                        | 9 |
| <i>Float f2;</i>                                        | 9 |
| <i>Size sz1;</i>                                        | 9 |
| <i>Size sz2;</i>                                        | 9 |
| <i>String str1;</i>                                     | 9 |
| <i>String str2;</i>                                     | 9 |
| <br>                                                    |   |
| <i>i1 + i2;</i>                                         | 9 |
| <i>f1 - i2;</i>                                         | 9 |
| <i>sz1 * sz2 * 2b;</i>                                  | 9 |
| <i>sz1 / 1024b;</i>                                     | 9 |
| <br>                                                    |   |
| <i>str1 = "basl";</i>                                   | 9 |
| <i>str2 = "cool";</i>                                   | 9 |
| <br>                                                    |   |
| <i>// the following is a string concatenation</i>       | 9 |
| <i>// operation resulting in the string:</i>            | 9 |
| <i>// "basl cool"</i>                                   | 9 |
| <i>str1 + str2;</i>                                     | 9 |
| <br>                                                    |   |
| <i>i1 % 10;</i>                                         | 9 |
| <br>                                                    |   |
| <i>i1 = 5;</i>                                          | 9 |
| <i>f1 = 7.5;</i>                                        | 9 |
| <br>                                                    |   |
| <i>str1 = "The values are: " + i1 + " and " + f1;</i>   | 9 |
| b) Unary expressions                                    | 9 |
| <br>                                                    |   |
| <b>+expr // positive - multiplies by 1 an</b>           | 9 |
| <b>// expression that is</b>                            | 9 |
| <b>// of Int, Float, or</b>                             | 9 |
| <b>// Size type</b>                                     | 9 |
| <br>                                                    |   |
| <b>-expr // negative - multiplies by -1 an</b>          | 9 |
| <b>// expression that is</b>                            | 9 |
| <b>// of Int, Float, or</b>                             | 9 |
| <b>// Size type</b>                                     | 9 |
| <br>                                                    |   |
| <b>!expr // not - converts a non-zero expr</b>          | 9 |
| <b>// value into 0, and a</b>                           | 9 |
| <b>// zero expr value into 1</b>                        | 9 |

|    |                                                                                                   |   |
|----|---------------------------------------------------------------------------------------------------|---|
|    | // <b>where expr type must be</b>                                                                 | 9 |
|    | // <b>of type Int or Float</b>                                                                    | 9 |
|    | Some sample unary expressions:                                                                    | 9 |
|    | <i>Int i;</i>                                                                                     | 9 |
|    | <i>+3;</i>                                                                                        | 9 |
|    | <i>-(i + 4);</i>                                                                                  | 9 |
|    | <i>!i;</i>                                                                                        | 9 |
| c) | Logical expressions                                                                               | 9 |
|    | <b>lexpr EQ rexpr</b>                                                                             | 9 |
|    | <b>lexpr NEQ rexpr</b>                                                                            | 9 |
|    | <b>lexpr LT rexpr</b>                                                                             | 9 |
|    | <b>lexpr LE rexpr</b>                                                                             | 9 |
|    | <b>lexpr GT rexpr</b>                                                                             | 9 |
|    | <b>lexpr GE rexpr</b>                                                                             | 9 |
|    | <b>lexpr AND rexpr</b>                                                                            | 9 |
|    | <b>lexpr OR rexpr</b>                                                                             | 9 |
|    | <b>lexpr</b> and <b>rexpr</b> must have types that are mutually consistent as shown               | 9 |
|    | in the following table:                                                                           | 9 |
|    | <i>lterminal-expr rterminal-expr</i>                                                              | 9 |
|    | =====                                                                                             | 9 |
|    | <i>Int or Float Int or Float</i>                                                                  | 9 |
|    | <i>Dayofweek Dayofweek</i>                                                                        | 9 |
|    | <i>DateTime DateTime</i>                                                                          | 9 |
|    | <i>String String</i>                                                                              | 9 |
|    | <i>Size Size</i>                                                                                  | 9 |
|    | <i>Server Server</i>                                                                              | 9 |
|    | <i>Que Que</i>                                                                                    | 9 |
|    | <i>Job Job</i>                                                                                    | 9 |
|    | <i>CNode CNode</i>                                                                                | 9 |
|    | <i>Set Server Set Server</i>                                                                      | 9 |
|    | <i>Set Que Set Que</i>                                                                            | 9 |
|    | <i>Set Job Set Job</i>                                                                            | 9 |
|    | <i>Set CNode Set CNode</i>                                                                        | 9 |
|    | For <b>AND</b> , <b>OR</b> operators, the <b>lexpr</b> , <b>rexpr</b> consistent types are Int or | 9 |
|    | Float.                                                                                            | 9 |
|    | Some sample logical expressions:                                                                  | 9 |
|    | <i>i1 EQ i2;</i>                                                                                  | 9 |
|    | <i>i1 NEQ f2;</i>                                                                                 | 9 |
|    | <i>dow1 LE dow2;</i>                                                                              | 9 |
|    | <i>d1 LT d2;</i>                                                                                  | 9 |
|    | <i>str1 GT str2;</i>                                                                              | 9 |
|    | <i>sz1 GE sz2;</i>                                                                                | 9 |
| d) | Post-operator expressions                                                                         | 9 |
|    | These are expressions that are merely shortcut to assignment state-                               | 9 |
|    | ments.                                                                                            | 9 |

|    |                                                                                                                                                                                                                                                                                   |   |
|----|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
|    | <b>IDENTIFIER++;</b> // <b>identifier=identifier+1</b>                                                                                                                                                                                                                            | 9 |
|    | <b>IDENTIFIER--;</b> // <b>identifier=identifier-1</b>                                                                                                                                                                                                                            | 9 |
|    | <b>IDENTIFIER</b> must be of Int or Float type.                                                                                                                                                                                                                                   | 9 |
|    | Example:                                                                                                                                                                                                                                                                          | 9 |
|    | <i>Int i;</i>                                                                                                                                                                                                                                                                     | 9 |
|    | <i>Float f;</i>                                                                                                                                                                                                                                                                   | 9 |
|    | <i>i++;</i>                                                                                                                                                                                                                                                                       | 9 |
|    | <i>f--;</i>                                                                                                                                                                                                                                                                       | 9 |
| e) | Function call                                                                                                                                                                                                                                                                     | 9 |
|    | <b>function-name ( arg1 ,arg2 ... , argN )</b>                                                                                                                                                                                                                                    | 9 |
|    | where <b>arg1, ..., argN</b> can be any constant or variable. You can't have another function call as an argument.                                                                                                                                                                | 9 |
|    | Example:                                                                                                                                                                                                                                                                          | 9 |
|    | <i>Void pr(Int a) {</i>                                                                                                                                                                                                                                                           | 9 |
|    | <i>print(a);</i>                                                                                                                                                                                                                                                                  | 9 |
|    | <i>}</i>                                                                                                                                                                                                                                                                          | 9 |
|    | <i>pr(5);</i>                                                                                                                                                                                                                                                                     | 9 |
|    | There are certain predefined functions that a scheduler writer can automatically call in his/her BASL code without a need to define it. These functions are referred to as assist functions (or helper functions) and they are discussed under <b>PREDEFINED FUNCTIONS</b> topic. | 9 |
| f) | Constants                                                                                                                                                                                                                                                                         | 9 |
|    | Some valid constant expressions are given in the following:                                                                                                                                                                                                                       | 9 |
|    | 5;                                                                                                                                                                                                                                                                                | 9 |
|    | +1.2;                                                                                                                                                                                                                                                                             | 9 |
|    | SUN;                                                                                                                                                                                                                                                                              | 9 |
|    | MON;                                                                                                                                                                                                                                                                              | 9 |
|    | TUE;                                                                                                                                                                                                                                                                              | 9 |
|    | WED;                                                                                                                                                                                                                                                                              | 9 |
|    | THU;                                                                                                                                                                                                                                                                              | 9 |
|    | FRI;                                                                                                                                                                                                                                                                              | 9 |
|    | SAT;                                                                                                                                                                                                                                                                              | 9 |
|    | (4   4   1997);                                                                                                                                                                                                                                                                   | 9 |
|    | (12:01:00);                                                                                                                                                                                                                                                                       | 9 |
|    | (4   4   1997@12:01:00);                                                                                                                                                                                                                                                          | 9 |
|    | "wonderful";                                                                                                                                                                                                                                                                      | 9 |
|    | -1b;                                                                                                                                                                                                                                                                              | 9 |
|    | SYNCRUN;                                                                                                                                                                                                                                                                          | 9 |
|    | ASYNCRUN;                                                                                                                                                                                                                                                                         | 9 |
|    | DELETE;                                                                                                                                                                                                                                                                           | 9 |
|    | RERUN;                                                                                                                                                                                                                                                                            | 9 |
|    | HOLD;                                                                                                                                                                                                                                                                             | 9 |
|    | RELEASE;                                                                                                                                                                                                                                                                          | 9 |
|    | SIGNAL;                                                                                                                                                                                                                                                                           | 9 |
|    | MODIFYATTR;                                                                                                                                                                                                                                                                       | 9 |
|    | MODIFYRES;                                                                                                                                                                                                                                                                        | 9 |
|    | (1, 3);                                                                                                                                                                                                                                                                           | 9 |



|                                                      |   |
|------------------------------------------------------|---|
| (2.3, 4.6);                                          | 9 |
| (WED, FRI);                                          | 9 |
| ((4   4   1997), (4   10   1997));                   | 9 |
| ((12:01:00), (12:30:00));                            | 9 |
| ((4   4   1997@12:01:00), (4   10   1997@12:30:00)); | 9 |
| (23gb, 50gb);                                        | 9 |
| NOSERVER;                                            | 9 |
| NOCNODE;                                             | 9 |
| NOQUE;                                               | 9 |
| NOJOB;                                               | 9 |
| EMPTYSETSERVER;                                      | 9 |
| EMPTYSETCNODE;                                       | 9 |
| EMPTYSETQUE;                                         | 9 |
| EMPTYSETJOB;                                         | 9 |
| NULLSTR;                                             | 9 |
| SUCCESS;                                             | 9 |
| FAIL;                                                | 9 |
| SERVER_ACTIVE;                                       | 9 |
| SERVER_IDLE;                                         | 9 |
| SERVER_SCHED;                                        | 9 |
| SERVER_TERM;                                         | 9 |
| SERVER_TERMDELAY;                                    | 9 |
| QTYPE_E;                                             | 9 |
| QTYPE_R;                                             | 9 |
| SCHED_DISABLED;                                      | 9 |
| SCHED_ENABLED;                                       | 9 |
| FALSE;                                               | 9 |
| TRUE;                                                | 9 |
| TRANSIT;                                             | 9 |
| QUEUED;                                              | 9 |
| HELD;                                                | 9 |
| WAITING;                                             | 9 |
| RUNNING;                                             | 9 |
| EXITING;                                             | 9 |
| CNODE_OFFLINE;                                       | 9 |
| CNODE_DOWN;                                          | 9 |
| CNODE_FREE;                                          | 9 |
| CNODE_RESERVE;                                       | 9 |
| CNODE_INUSE_EXCLUSIVE;                               | 9 |
| CNODE_INUSE_SHARED;                                  | 9 |
| CNODE_TIMESHARED;                                    | 9 |
| CNODE_CLUSTER;                                       | 9 |
| CNODE_UNKNOWN;                                       | 9 |
| OP_EQ;                                               | 9 |
| OP_NEQ;                                              | 9 |
| OP_LE;                                               | 9 |
| OP_LT;                                               | 9 |
| OP_GE;                                               | 9 |
| OP_GT;                                               | 9 |
| OP_MAX;                                              | 9 |
| OP_MIN;                                              | 9 |
| g) Identifier                                        | 9 |
| Example:                                             | 9 |

|                                                                                                                                                                                                                                              |   |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| Int i;                                                                                                                                                                                                                                       | 9 |
| i;                                                                                                                                                                                                                                           | 9 |
| 2. Assignment statement                                                                                                                                                                                                                      | 9 |
| <b>IDENTIFIER = expr ;</b>                                                                                                                                                                                                                   | 9 |
| <b>IDENTIFIER</b> and <b>expr</b> must have types that are mutually consistent as illustrated in the following table:                                                                                                                        | 9 |
| <i>identifier</i> <i>expr</i>                                                                                                                                                                                                                | 9 |
| =====                                                                                                                                                                                                                                        | 9 |
| Int          Int, Float                                                                                                                                                                                                                      | 9 |
| Float        Int, Float                                                                                                                                                                                                                      | 9 |
| Dayofweek    Dayofweek                                                                                                                                                                                                                       | 9 |
| DateTime     DateTime                                                                                                                                                                                                                        | 9 |
| String        String, Int, Float, Dayofweek, Size                                                                                                                                                                                            | 9 |
| Size         Size                                                                                                                                                                                                                            | 9 |
| Que          Que                                                                                                                                                                                                                             | 9 |
| Job          Job                                                                                                                                                                                                                             | 9 |
| CNode        CNode                                                                                                                                                                                                                           | 9 |
| Server        Server                                                                                                                                                                                                                         | 9 |
| Dayofweek    Dayofweek                                                                                                                                                                                                                       | 9 |
| DateTime     DateTime                                                                                                                                                                                                                        | 9 |
| Set Server    Set Server                                                                                                                                                                                                                     | 9 |
| Set Que      Set Que                                                                                                                                                                                                                         | 9 |
| Set Job      Set Job                                                                                                                                                                                                                         | 9 |
| Set CNode    Set CNode                                                                                                                                                                                                                       | 9 |
| Range Int    Range Int                                                                                                                                                                                                                       | 9 |
| Range Float  Range Float                                                                                                                                                                                                                     | 9 |
| Range Dayofweek  Range Dayofweek                                                                                                                                                                                                             | 9 |
| Range DateTime  Range DateTime                                                                                                                                                                                                               | 9 |
| Range Size    Range Size                                                                                                                                                                                                                     | 9 |
| 3. if...else statement                                                                                                                                                                                                                       | 9 |
| The format of an if statement is similar to that in C with the delimiting "{" and "}" always present:                                                                                                                                        | 9 |
| <b>if( expr ) {</b>                                                                                                                                                                                                                          | 9 |
| <b>zero or more (true) STATEMENTS</b>                                                                                                                                                                                                        | 9 |
| <b>}</b>                                                                                                                                                                                                                                     | 9 |
| <b>if( expr ) {</b>                                                                                                                                                                                                                          | 9 |
| <b>zero or more (true) STATEMENTS</b>                                                                                                                                                                                                        | 9 |
| <b>} else {</b>                                                                                                                                                                                                                              | 9 |
| <b>zero or more (false) STATEMENTS</b>                                                                                                                                                                                                       | 9 |
| <b>}</b>                                                                                                                                                                                                                                     | 9 |
| The <b>expr</b> 's type must be either Int or Float, and after evaluation if its value is non-zero, then the true statements are executed. On the second form, if the <b>expr</b> evaluates to zero, then the false statements are executed. | 9 |
| Some sample <b>if</b> statements are given below:                                                                                                                                                                                            | 9 |
| <i>if( 2 * x )</i>                                                                                                                                                                                                                           | 9 |

```

{
    y = y + 3;
    print(y);
}

```

```

if (2 * x) {
    y = y + 3;
} else {
    if (3 * x) {
        y = 4;
    } else {
        y = 5;
    }
}

```

#### 4. for loop statement

The format of a for statement is as follows:

```

for( start; test; action ) {
    zero or more STATEMENTS
}

```

Just like in C, **for** first executes **start**, then evaluates the **test** condition to see if it returns a non-zero value. If it does, the **for** statements are executed. After the **for** statements are executed, then **action** is evaluated, and then it checks the **test** condition again in the same manner as before. **start** and **action** can be a simple assignment expression or a post-operator expression. **test** is a logical/relational expression. Some sample for statements are given in the following:

```

for (i = 0; i LT 3; i = i + 1)
{
    print(i);
}

```

```

for (i = 0; i LT 2 * x; i++)
{
    if (x GT 3)
    {
        y = 99;
    } else
    {
        x = 73;
    }
}

```

#### 5. foreach loop statement

This statement is primarily used for successively retrieving each element of a Set data type: Set Server, Set CNode, Set Job, Set Que. The syntax is:

```

foreach ( IDENTIFIER1 in IDENTIFIER2 ) {
    zero or more STATEMENTS
}

```

where the following pairing of types for the identifiers are allowed: 9

```
IDENTIFIER1 IDENTIFIER2 9
=====
Server      Set Server 9
Que         Set Que    9
Job         Set Job    9
CNode      Set CNode   9
```

Example: 9

```
Server s; 9
Que q; 9
Job j; 9
CNode c; 9
```

```
Set Server ss; 9
Set Que sq; 9
Set Job sj; 9
Set CNode sc; 9
```

```
foreach(s in ss){ 9
    print(s); 9
} 9
foreach(q in sq){ 9
    print(q); 9
} 9
foreach(j in sj){ 9
    print(j); 9
} 9
foreach(c in sc){ 9
    print(c); 9
} 9
```

6. while loop statement 9

The syntax of a while loop is: 9

```
while ( expr ) { 9
    zero or more STATEMENTS 9
} 9
```

where **expr** must be of Int or Float type. If **expr** is non-zero, then the zero or more *STATEMENTS* are executed and **expr** is re-evaluated. 9

Example: 9

```
Int i; 9
i = 3; 9
while(i) { 9
    if( i EQ 0 ) { 9
        print("break on i = 1"); 9
        break; 9
    } 9
    i--; 9
} 9
```

7. switch statement 9  
 The switch statement is a multi-way decision that tests whether an identifier's value matches one of a number of values, and branches to a group of statements accordingly. 9  
 The syntax for a switch statement is: 9

```

switch( IDENTIFIER ) { 9
  case constant-expr : 9
  { 9
    zero or more STATEMENTS 9
  } 9
  case constant-expr : 9
  { 9
    zero or more STATEMENTS 9
  } 9
  ... 9
  case in constant-rangeOrSet-expr : 9
  { 9
    zero or more STATEMENTS 9
  } 9
  case in IDENTIFIER-rangeOrSettype : 9
  { 9
    zero or more STATEMENTS 9
  } 9
  default : 9
  { 9
    zero or more STATEMENTS 9
  } 9
} 9

```

where **constant-expr** is an **expr** of type Int, Float, Dayofweek, DateTime, Size, String, Server, Que, Job, or CNode. **constant-rangeOrSet-expr** and **IDENTIFIER-rangeOrSettype** can be of type Set Server, Set CNode, Set Que, Set Job, Range Int, Range Float, Range Dayofweek, Range DateTime, or Range Size. 9

**IDENTIFIER** cannot be of type Void. **IDENTIFIER**'s type must be consistent with **constant-expr**'s, **constant-rangeOrSet-expr**'s, and **IDENTIFIER-rangeOrSettype**'s type as illustrated in the following table: 9

|                                                             |   |
|-------------------------------------------------------------|---|
| <i>IDENTIFIER constant-range-expr, IDENTIFIER-rangetype</i> | 9 |
| =====                                                       | 9 |
| <i>Server Set Server</i>                                    | 9 |
| <i>Que Set Que</i>                                          | 9 |
| <i>Job Set Job</i>                                          | 9 |
| <i>CNode Set CNode</i>                                      | 9 |
| <i>Int Range Int</i>                                        | 9 |
| <i>Float Range Float</i>                                    | 9 |
| <i>Dayofweek Range Dayofweek</i>                            | 9 |
| <i>DateTime Range DateTime</i>                              | 9 |
| <i>Size Range Size</i>                                      | 9 |

If a case expression matches the **IDENTIFIER**'s value, then the corresponding block of statements are executed. Unlike in C, execution does NOT fall through to the next case statement. The reason for this is that *basl2c* will translate this **switch** statement into if-elseif-else construct. The case labeled default is executed if none of the other cases are satisfied. The **default** is optional; if it isn't there, and if none of the cases match, no action takes place.

Example: 9  
*Dayofweek dow;* 9

```

switch(dow) 9
{ 9
    case MON: 9
    { 9
        print("case MON"); 9
    } 9
    case TUE: 9
    { 9
        print("case TUE"); 9
    } 9
    case WED: 9
    { 9
        print("case WED"); 9
    } 9
    case THU: 9
    { 9
        print("case THU"); 9
    } 9
    case FRI: 9
    { 9
        print("case FRI"); 9
    } 9
    case SAT: 9
    { 9
        print("case SAT"); 9
    } 9
    case SUN: 9
    { 9
        print("case SUN"); 9
    } 9
    default: 9
    { 9

```

|                                                                                                                                                                             |   |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| <i>print("case defaulted");</i>                                                                                                                                             | 9 |
| }                                                                                                                                                                           | 9 |
| }                                                                                                                                                                           | 9 |
| <br><i>Int a;</i>                                                                                                                                                           | 9 |
| <i>Range Int ri;</i>                                                                                                                                                        | 9 |
| <i>ri = (10, 12);</i>                                                                                                                                                       | 9 |
| <i>switch(a)</i>                                                                                                                                                            | 9 |
| {                                                                                                                                                                           | 9 |
| <i>case in (1,5):</i>                                                                                                                                                       | 9 |
| {                                                                                                                                                                           | 9 |
| <i>print("case 1,5");</i>                                                                                                                                                   | 9 |
| }                                                                                                                                                                           | 9 |
| <i>case in (6,9):</i>                                                                                                                                                       | 9 |
| {                                                                                                                                                                           | 9 |
| <i>print("case 6,9");</i>                                                                                                                                                   | 9 |
| }                                                                                                                                                                           | 9 |
| <i>case in ri:</i>                                                                                                                                                          | 9 |
| {                                                                                                                                                                           | 9 |
| <i>print("case ri");</i>                                                                                                                                                    | 9 |
| }                                                                                                                                                                           | 9 |
| }                                                                                                                                                                           | 9 |
| 8. print statement                                                                                                                                                          | 9 |
| Print statement is capable of printing to stdout the value of any <b>identifier</b> or <b>constant</b> of type Int, Float, Dayofweek, DateTime, String, Size, Que, Job, CN- | 9 |
| ode, Server, Range Int, Range Float, Range Dayofweek, Range DateTime,                                                                                                       | 9 |
| Range Size.                                                                                                                                                                 | 9 |
| The syntax is as follows:                                                                                                                                                   | 9 |
| <br><b>print ( IDENTIFIER );</b>                                                                                                                                            | 9 |
| <b>print ( constant );</b>                                                                                                                                                  | 9 |
| <br>Example:                                                                                                                                                                | 9 |
| <i>DateTime dt;</i>                                                                                                                                                         | 9 |
| <i>CNode cn;</i>                                                                                                                                                            | 9 |
| <br><i>dt = (4/4/1997@12:13:36);</i>                                                                                                                                        | 9 |
| <i>cn = AllNodesLocalHostGet();</i>                                                                                                                                         | 9 |
| <br><i>print(dt);</i>                                                                                                                                                       | 9 |
| <i>print(cn);</i>                                                                                                                                                           | 9 |
| <br>For Set types, use <b>foreach</b> to go through each element and print as in:                                                                                           | 9 |
| <br><i>Server s;</i>                                                                                                                                                        | 9 |
| <i>Set Server ss;</i>                                                                                                                                                       | 9 |
| <br><i>ss = AllServersGet();</i>                                                                                                                                            | 9 |
| <br><i>foreach(s in ss) {</i>                                                                                                                                               | 9 |
| <i>print(s);</i>                                                                                                                                                            | 9 |
| }                                                                                                                                                                           | 9 |

9. continue statement 9

**continue ;** 9

The **continue** statement must have been invoked within a **for**, **foreach**, and **while** loop. It causes the next iteration of the enclosing loop to begin. 9

10. break statement 9

**break ;** 9

The **break** statement must have been invoked within a **for**, **foreach**, and **while** loop. It provides an early exit from the enclosing loop. 9

11. return statement 9

**return(IDENTIFIER) ;** 9

**return(constant) ;** 9

**return() ;** 9

The return statement provides the value (if any) to be returned by a function. The type returned by **IDENTIFIER** and **constant** must match the calling function's return type. **constant** types allowed are anything except Set and Range types. The last format, **return()** is usually called within a function that doesn't return any value ( like **sched\_main()** ). 9

12. exit statement 9

**exit(constant);** 9

where **constant** is of type Int. Calling this will terminate the scheduler. 9

13. Comment statement 9

These are statements prefixed by `"/"` and they are ignored by the BASL compiler. 9

*// this line is ignored* 9

*Int i; // string following the slashes is ignored* 9

**OPERATOR PRECEDENCE AND ASSOCIATIVITY** 9

The following table shows the various operator precedence levels and associativity defined in the BASL language. The operators are listed in the order of decreasing precedence. The higher the precedence of an operator, the earlier it gets executed. The order in which the operators on the same level are executed depends on the associativity: left means the operators are seen from left to right, while right means they are seen from right to left. 9

| <i>Operator</i>                               | <i>Associativity</i> | 9 |
|-----------------------------------------------|----------------------|---|
|                                               |                      |   |
| <i>! ++ -- + (unary plus) - (unary minus)</i> | <i>right</i>         | 9 |
| <i>* / %</i>                                  | <i>left</i>          | 9 |
| <i>+ -</i>                                    | <i>left</i>          | 9 |
| <i>LT LE GT GE</i>                            | <i>left</i>          | 9 |



|               |              |   |
|---------------|--------------|---|
| <i>EQ NEQ</i> | <i>left</i>  | 9 |
| <i>AND</i>    | <i>left</i>  | 9 |
| <i>OR</i>     | <i>left</i>  | 9 |
| <i>=</i>      | <i>right</i> | 9 |

## PREDEFINED FUNCTIONS 9

In BASL(2), a **Server** data type maps directly to a batch server object. Similarly, **CNode** is to mom/resmom, **Job** is to batch job, and **Que** is to batch queue. However, not all attributes to the PBS objects can be accessed from BASL. Only a subset of attributes, those that seemed to make sense in the context of a scheduler, are made available, and values to these attributes can be accessed by calling the following pre-defined functions, known also as assist/helper functions. 9

### (1) Server-related functions 9

#### **Set Server AllServersGet(void) 9**

Returns the list of servers specified in the configuration file for which the scheduler writer wants the system to periodically check for status, queues and jobs info. See **pbs\_sched\_basl(8B)** for a discussion on the format of the configuration file. 9

CAUTION: This function must be called from inside **sched\_main()** so that at every scheduling iteration, the most up to date **Set Server** structure is returned. 9

#### **Server AllServersLocalHostGet(void) 9**

Returns the Server object that represents the local host. unset value: NOSERVER. This is a simple function to call for non-cluster environments where only one server host exists. 9

CAUTION: This function must be called from inside **sched\_main()** (or from within function called by **sched\_main**) so that at every scheduling iteration, the most up to date **Server** structure is returned. 9

#### **String ServerInetAddrGet(Server s) 9**

Returns name for Server s. unset value: NULLSTR 9

#### **String ServerDefQueGet(Server s) 9**

Returns the default\_queue attribute of Server s. unset value: NULLSTR 9

#### **Int ServerStateGet(Server s) 9**

Returns server\_state attribute of Server s. 9

Return value: 9

SERVER\_ACTIVE, SERVER\_IDLE, SERVER\_SCHED, SERV-  
ER\_TERM, SERVER\_TERMDelay, -1 (unset value) 9

#### **Int ServerMaxRunJobsGet(Server s) 9**

Returns max\_running attribute of Server s. unset value: 0 9

#### **Int ServerMaxRunJobsPerUserGet(Server s) 9**

Returns max\_user\_run attribute of Server s. unset value: 0 9

#### **Int ServerMaxRunJobsPerGroupGet(Server s) 9**

Returns max\_group\_run attribute of Server s. unset value: 0 9

#### **Set Que ServerQueuesGet(Server s) 9**

Returns list of queues managed by Server s. 9

#### **Set Job ServerJobsGet(Server s) 9**

Returns list of jobs managed by Server s. For obtaining a subset of this list, see **QueJobsGet()**. 9

#### **Int ServerIntResAvailGet(Server s, String name) 9**

Returns the value to resource specified in **name** that is available to jobs run by this server (Server resources\_available.name attribute). Call this func- 9

tion for resources with values that are of Int type. Sample resource names are: cput, pcp, walltime, mppt, pmppt, nice, procs, mppe, ncpus, pncpus, nodect, srfs\_assist, mta,..., mth. For a description of these resource names, see pbs\_resources\_irix5(7B), pbs\_resources\_sp2(7B), pbs\_resources\_sunos4(7B), pbs\_resources\_unicos8(7B), pbs\_server\_attributes(7B), pbs\_resources\_irix6(7B), pbs\_resources\_linux(7B).

Example:  
*Int cpuAvail;*  
 // return the # of cpus currently available in  
 // the server  
*cpuAvail = ServerIntResAvailGet(server, "ncpus");*

#### **Size ServerSizeResAvailGet(Server s, String name)**

Returns the value to resource specified in **name** that is available to jobs run by this server (Server resources\_available.name attribute). Call this function for resources with values that are of Size type. Sample resource names are: file, mem, pmem, workingset, pf, ppf, srfs\_tmp, srfs\_wrk, srfs\_big, srfs\_fast, sds, psds. For a description of these resource names, see pbs\_resources\_irix5(7B), pbs\_resources\_sp2(7B), pbs\_resources\_sunos4(7B), pbs\_resources\_unicos8(7B), pbs\_server\_attributes(7B), pbs\_resources\_irix6(7B), pbs\_resources\_linux(7B).

Example:  
*Size memAvail;*  
 // return the amount of available memory in  
 // the server  
*memAvail = ServerSizeResAvailGet(server, "mem");*

#### **String ServerStringResAvailGet(Server s, String name)**

Returns the value to resource specified in **name** that is available to jobs run by this server (Server resources\_available.name attribute). Call this function for resources with values that are of String type. Sample resource names are: nodes, arch, neednodes. For a description of these resource names, see pbs\_resources\_irix5(7B), pbs\_resources\_sp2(7B), pbs\_resources\_sunos4(7B), pbs\_resources\_unicos8(7B), pbs\_server\_attributes(7B), pbs\_resources\_irix6(7B), pbs\_resources\_linux(7B).

Example:  
*String type;*  
 // return the architecture (or os type) of  
 // the server  
*type = ServerStringResAvailGet(server, "arch");*

#### **Int ServerIntResAssignGet(Server s, String name)**

Returns the value to resource specified in **name** that is allocated to running jobs (Server resources\_assigned.name attribute). Call this function for resources with values that are of Int type. Sample resource names are: cput, pcp, walltime, mppt, pmppt, nice, procs, mppe, ncpus, pncpus, nodect, srfs\_assist, mta,..., mth. For a description of these resource names, see pbs\_resources\_irix5(7B), pbs\_resources\_sp2(7B), pbs\_resources\_sunos4(7B), pbs\_resources\_unicos8(7B), pbs\_server\_attributes(7B), pbs\_resources\_irix6(7B), pbs\_resources\_linux(7B).

Example:  
*Int cpuAssn;*  
 // return the # of cpus currently assigned in

```

// the server
cpuAssn = ServerIntResAssignGet(server, "ncpus");

```

**Size ServerSizeResAssignGet(Server s, String name)**

Returns the value to resource specified in **name** that is allocated to running jobs (Server resources\_assigned.name attribute). Call this function for resources with values that are of Size type. Sample resource names are: file, mem, pmem, workingset, pf, ppf, srfs\_tmp, srfs\_wrk, srfs\_big, srfs\_fast, sds, psds. For a description of these resource names, see pbs\_resources\_irix5(7B), pbs\_resources\_sp2(7B), pbs\_resources\_sunos4(7B), pbs\_resources\_unicos8(7B), pbs\_server\_attributes(7B), pbs\_resources\_irix6(7B), pbs\_resources\_linux(7B).

Example:

```

Size sdsAssn;
// return the amount of sds space currently assigned
// in the server
sdsAssn = ServerSizeResAssignGet(server, "sds");

```

**String ServerStringResAssignGet(Server s, String name)**

Returns the value to resource specified in **name** that is allocated to running jobs (Server resources\_assigned.name attribute). Call this function for resources with values that are of String type. Sample resource names are: nodes, arch, neednodes. For a description of these resource names, see pbs\_resources\_irix5(7B), pbs\_resources\_sp2(7B), pbs\_resources\_sunos4(7B), pbs\_resources\_unicos8(7B), pbs\_server\_attributes(7B), pbs\_resources\_irix6(7B), pbs\_resources\_linux(7B).

**Set CNode ServerNodesGet(Server s)**

Returns the set of nodes managed by server s. unset value: EMPTYSETCNODE.

NOTE: You can usually call the following functions for the nodes returned by this call: CNodeStateGet(), CNodePropertiesGet(), and CNodeTypeGet().

**Int ServerNodesQuery(Server s, String spec)**

Issues a request to the specified server to query the availability of resources specified in **spec**. At the present time, the only resource specification allowed is one that involves "nodes" and it can be of the format "nodes", "nodes=", or "nodes=<type>". The query results can be accessed by calling the following functions: ServerNodesNumAvailGet(), ServerNodesNumAllocGet(), ServerNodesNumRsvdGet(), ServerNodesNumDownGet().

NOTE: This is a wrapper to the pbs\_resquery(3B) server function.

Return value:  
SUCCESS, FAIL

**Int ServerNodesNumAvailGet(Server s)**

Returns the number of nodes available for those managed by the specified server, or as reflected by the most recent query specified by ServerNodesQuery(). If the return value is zero, then this means that some number of nodes currently needed to satisfy the specification of ServerNodesQuery() are currently unavailable. The request maybe satisfied at some later time. If the result is negative, no combination of known nodes can satisfy the specification.

**Int ServerNodesNumAllocGet(Server s)**

Returns the number of nodes allocated for those managed by the specified server, or as reflected by the most recent query specified by ServerNodesQuery().

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |   |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| <b>Int ServerNodesNumRsvdGet(Server s)</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | 9 |
| Returns the number of nodes reserved for those managed by the specified server, or as reflected by the most recent query specified by ServerNodesQuery().                                                                                                                                                                                                                                                                                                                                                                                                                                                      | 9 |
| <b>Int ServerNodesNumDownGet(Server s)</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | 9 |
| Returns the number of nodes down for those managed by the specified server, or as reflected by the most recent query specified by ServerNodesQuery().                                                                                                                                                                                                                                                                                                                                                                                                                                                          | 9 |
| <b>Int ServerNodesReserve(Server s,String spec,Int resId)</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | 9 |
| Issues a request to the specified server to reserve the resources specified in <b>spec</b> . A value of 0 for <b>resId</b> means that this is for doing a new reservation. Otherwise, the number will represent an existing (partial) reservation. Resources currently reserved for this <b>resId</b> will be released and the full reservation will be attempted again. At the present time the only resources which may be specified are "nodes". It should be specified as <b>nodes=specification</b> where specification is what a user specifies in the -l option argument list for nodes, see qsub (1B). | 9 |
| NOTE: This is a wrapper to the pbs_resreserve(3B) server function.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | 9 |
| Return value:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | 9 |
| a reference number to a successful or partially-successful reservation, or FAIL                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | 9 |
| <b>Int ServerNodesRelease(Server s, Int resId)</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | 9 |
| This releases or frees resources reserved with the reference number specified in <b>resId</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | 9 |
| NOTE: This is a wrapper to the pbs_resrelease(3B) server function.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | 9 |
| Return value:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | 9 |
| SUCCESS, or FAIL                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | 9 |
| (2) Que-related functions:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | 9 |
| <b>String QueNameGet( Que que )</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | 9 |
| Returns name of Que que. unset value: NULLSTR                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | 9 |
| <b>Int QueTypeGet( Que que )</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | 9 |
| Returns queue_type attribute of Que que.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | 9 |
| Return value: QTYPE_E (Execution), QTYPE_R (Routing), -1 (unset value)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | 9 |
| <b>Int QueNumJobsGet( Que que )</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | 9 |
| Returns number of jobs residing in Que que. unset value: 0                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | 9 |
| <b>Int QueMaxRunJobsGet( Que que )</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | 9 |
| Returns max_running attribute of Que que. unset value: 0                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | 9 |
| <b>Int QueMaxRunJobsPerUserGet( Que que )</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | 9 |
| Returns max_user_run attribute of Que que. unset value: 0                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | 9 |
| <b>Int QueMaxRunJobsPerGroupGet( Que que )</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | 9 |
| Returns max_group_run attribute of Que que. unset value: 0                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | 9 |
| <b>Int QuePriorityGet( Que que )</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | 9 |
| Returns Priority attribute of Que que. unset value: 0                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | 9 |
| <b>Int QueStateGet( Que que )</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | 9 |
| Returns started attribute of Que que - the job execution selection state of the que: SCHED_DISABLED, SCHED_ENABLED. unset value: SCHED_DISABLED                                                                                                                                                                                                                                                                                                                                                                                                                                                                | 9 |
| <b>Set Job QueJobsGet( Que que )</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | 9 |
| Returns the list of jobs currently residing in que.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | 9 |

- Int QueIntResAvailGet(Que q, String name)** 9  
 Returns the value to resource specified in **name** that is available to jobs 9  
 running from this q (Que resources\_available.name attribute). Call this func- 9  
 tion for resources with values that are of Int type. Sample resource 9  
 names are: cput, pcpur, walltime, mppt, pmppt, nice, procs, mppe, ncpus, 9  
 pncpus, nodect, srfs\_assist, mta,..., mth. For a description of these re- 9  
 source names, see pbs\_resources\_irix5(7B), pbs\_resources\_sp2(7B), 9  
 pbs\_resources\_sunos4(7B), pbs\_resources\_unicos8(7B), pbs\_server\_at- 9  
 tributes(7B), pbs\_resources\_irix6(7B), pbs\_resources\_linux(7B). 9
- Size QueSizeResAvailGet(Que q, String name)** 9  
 Returns the value to resource specified in **name** that is available to jobs 9  
 running from this q (Que resources\_available.name attribute). Call this func- 9  
 tion for resources with values that are of Size type. Sample resource 9  
 names are: file, mem, pmem, workingset, pf, ppf, srfs\_tmp, srfs\_wrk, 9  
 srfs\_big, srfs\_fast, sds, psds. For a description of these resource names, 9  
 see pbs\_resources\_irix5(7B), pbs\_resources\_sp2(7B), pbs\_re- 9  
 sources\_sunos4(7B), pbs\_resources\_unicos8(7B), pbs\_server\_at- 9  
 tributes(7B), pbs\_resources\_irix6(7B), pbs\_resources\_linux(7B). 9
- String QueStringResAvailGet(Que q, String name)** 9  
 Returns the value to resource specified in **name** that is available to jobs 9  
 running from this q (Que resources\_available.name attribute). Call this func- 9  
 tion for resources with values that are of String type. Sample resource 9  
 names are: nodes, arch, neednodes. For a description of these resource 9  
 names, see pbs\_resources\_irix5(7B), pbs\_resources\_sp2(7B), pbs\_re- 9  
 sources\_sunos4(7B), pbs\_resources\_unicos8(7B), pbs\_server\_at- 9  
 tributes(7B), pbs\_resources\_irix6(7B), pbs\_resources\_linux(7B). 9
- Int QueIntResAssignGet(Que q, String name)** 9  
 Returns the value to resource specified in **name** that is allocated to jobs 9  
 running from this queue (Que resources\_assigned.name attribute). Call this 9  
 function for resources with values that are of Int type. Sample resource 9  
 names are: cput, pcpur, walltime, mppt, pmppt, nice, procs, mppe, ncpus, 9  
 pncpus, nodect, srfs\_assist, mta,..., mth. For a description of these re- 9  
 source names, see pbs\_resources\_irix5(7B), pbs\_resources\_sp2(7B), 9  
 pbs\_resources\_sunos4(7B), pbs\_resources\_unicos8(7B), pbs\_server\_at- 9  
 tributes(7B), pbs\_resources\_irix6(7B), pbs\_resources\_linux(7B). 9
- Size QueSizeResAssignGet(Que q, String name)** 9  
 Returns the value to resource specified in **name** that is allocated to jobs 9  
 running from this q (Que resources\_assigned.name attribute). Call this 9  
 function for resources with values that are of Size type. Sample resource 9  
 names are: file, mem, pmem, workingset, pf, ppf, srfs\_tmp, srfs\_wrk, 9  
 srfs\_big, srfs\_fast, sds, psds. For a description of these resource names, 9  
 see pbs\_resources\_irix5(7B), pbs\_resources\_sp2(7B), pbs\_re- 9  
 sources\_sunos4(7B), pbs\_resources\_unicos8(7B), pbs\_server\_at- 9  
 tributes(7B), pbs\_resources\_irix6(7B), pbs\_resources\_linux(7B). 9
- String QueStringResAssignGet(Que q, String name)** 9  
 Returns the value to resource specified in **name** that is allocated to jobs 9  
 running from this q (Que resources\_assigned.name attribute). Call this 9  
 function for resources with values that are of String type. Sample re- 9  
 source names are: nodes, arch, neednodes. For a description of these re- 9  
 source names, see pbs\_resources\_irix5(7B), pbs\_resources\_sp2(7B), 9  
 pbs\_resources\_sunos4(7B), pbs\_resources\_unicos8(7B), pbs\_server\_at- 9  
 tributes(7B), pbs\_resources\_irix6(7B), pbs\_resources\_linux(7B). 9

|                                                                             |   |
|-----------------------------------------------------------------------------|---|
| (3) Job-related functions                                                   | 9 |
| <b>String JobIdGet( Job job )</b>                                           | 9 |
| Returns job identifier of Job job. unset value: NULLSTR                     | 9 |
| <b>String JobNameGet( Job job )</b>                                         | 9 |
| Returns Job_Name attribute of Job job. unset value: NULLSTR                 | 9 |
| <b>String JobOwnerNameGet( Job job )</b>                                    | 9 |
| Returns Job_Owner attribute of Job job. unset value: NULLSTR                | 9 |
| <b>String JobEffectiveUserNameGet( Job job )</b>                            | 9 |
| Returns euser attribute of Job job.                                         | 9 |
| <b>String JobEffectiveGroupNameGet( Job job )</b>                           | 9 |
| Returns egroup attribute of Job job. unset value: NULLSTR                   | 9 |
| <b>Int JobStateGet ( Job job )</b>                                          | 9 |
| Returns job_state attribute of Job job.                                     | 9 |
| Return value:                                                               | 9 |
| TRANSIT, QUEUED, HELD, WAITING, RUNNING, EXITING, -1                        | 9 |
| (unset value)                                                               | 9 |
| <b>Int JobPriorityGet( Job job )</b>                                        | 9 |
| Returns Priority attribute of Job job. unset value: 0                       | 9 |
| <b>Int JobRerunFlagGet( Job job )</b>                                       | 9 |
| Returns Rerunable attribute of Job job.                                     | 9 |
| Return value: FALSE, TRUE, -1 (unset value)                                 | 9 |
| <b>Int JobInteractiveFlagGet( Job job )</b>                                 | 9 |
| Returns interactive attribute of Job job.                                   | 9 |
| Return value: FALSE, TRUE. unset value: FALSE                               | 9 |
| <b>DateTime JobDateTimeCreatedGet( Job job )</b>                            | 9 |
| Returns the ctime attribute of Job job. unset value: (0 0 0@-1:-1:-1)       | 9 |
| <b>String JobEmailAddrGet( Job job )</b>                                    | 9 |
| Returns the Mail_Users attribute of Job job. unset value: NULLSTR           | 9 |
| <b>String JobStageinFilesGet( Job job )</b>                                 | 9 |
| Returns the stagein attribute of Job job. unset value: NULLSTR              | 9 |
| <b>String JobStageoutFilesGet( Job job )</b>                                | 9 |
| Returns stageout attribute of Job job. unset value: NULLSTR                 | 9 |
| <b>Set CNode JobNodesGet( Job job )</b>                                     | 9 |
| Returns the set of nodes (execution hosts) currently assigned to 'job'. un- | 9 |
| set value: EMPTYSETCNODE.                                                   | 9 |
| Example: Given a sample job j has exec_host set to "aspasia/0+alaba-        | 9 |
| ma/0":                                                                      | 9 |
| <i>jobnodes = JobNodesGet(j);</i>                                           | 9 |
| <i>foreach(n in jobnodes) {</i>                                             | 9 |
| <i>print("NODES ASSIGNED TO JOB");</i>                                      | 9 |
| <i>print(n);</i>                                                            | 9 |
| <i>processors_assn = CNodeNumCpusGet(n);</i>                                | 9 |
| <i>}</i>                                                                    | 9 |
| ---                                                                         | 9 |
| <i>will give the output:</i>                                                | 9 |
| <i>NODES ASSIGNED TO JOB:</i>                                               | 9 |

```

Node Name = aspasia 9
Properties = aspasia/0 9
# of Processors = 1 9
... 9
Node Name = alabama 9
Properties = alabama/0 9
# of Processors = 1 9
... 9
NOTE: CNodeNumCpusGet(n) will return 9
the # of cpus assigned for that node. 9
The output of this call will always 9
reflect the recent JobAction() 9
of the scheduler (always up to date). 9
Int JobIntResReqGet(Job job, String name) 9
Returns the value to resource specified in name as required by the job 9
(Job Resource_List.name attribute). Call this function for resources with 9
values that are of Int type. Sample resource names are: cput, pcpur, wall- 9
time, mppt, pmppt, nice, procs, mppe, ncpus, pncpus, nodect, srfs_assist, 9
mta,..., mth. For a description of these resource names, see pbs_re- 9
sources_irix5(7B), pbs_resources_sp2(7B), pbs_resources_sunos4(7B), 9
pbs_resources_unicos8(7B), pbs_server_attributes(7B), pbs_resources_ir- 9
ix6(7B), pbs_resources_linux(7B). 9

Example: 9
Int cputReq; 9
// returns the cput requirement of the job 9
cputReq = JobIntResReqGet(job, "cput"); 9

Size JobSizeResReqGet(Job job, String name) 9
Returns the value to resource specified in name as required by the job 9
(Job Resource_List.name attribute). Call this function for resources with 9
values that are of Size type. Sample resource names are: file, mem, 9
pmem, workingset, pf, ppf, srfs_tmp, srfs_wrk, srfs_big, srfs_fast, sds, ps- 9
ds. For a description of these resource names, see pbs_resources_ir- 9
ix5(7B), pbs_resources_sp2(7B), pbs_resources_sunos4(7B), pbs_re- 9
sources_unicos8(7B), pbs_server_attributes(7B), pbs_resources_irix6(7B), 9
pbs_resources_linux(7B). 9

Example: 9
Size memReq; 9
// returns the memory requirement of the job 9
memReq = JobSizeResReqGet(job, "mem"); 9

String JobStringResReqGet(Job job, String name) 9
Returns the value to resource specified in name as required by the job 9
(Job Resource_List.name attribute). Call this function for resources with 9
values that are of String type. Sample resource names are: nodes, arch, 9
neednodes. For a description of these resource names, see pbs_re- 9
sources_irix5(7B), pbs_resources_sp2(7B), pbs_resources_sunos4(7B), 9
pbs_resources_unicos8(7B), pbs_server_attributes(7B), pbs_resources_ir- 9
ix6(7B), pbs_resources_linux(7B). 9

Example: 9
String nodes; 9
// returns the nodes requirement property of 9
// the job 9

```

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |                                           |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------|
| <i>nodes = JobStringResReqGet(job, "nodes");</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | 9                                         |
| <b>Int JobIntResUseGet(Job job, String name)</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | 9                                         |
| Returns the value to resource specified in <b>name</b> used by the job (Job resources_used.name attribute). Call this function for resources with values that are of Int type. Sample resource names are: cput, pcpur, walltime, mppt, pmppt, nice, procs, mppe, ncpus, pncpus, nolect, srfs_assist, mta,..., mth. For a description of these resource names, see pbs_resources_irix5(7B), pbs_resources_sp2(7B), pbs_resources_sunos4(7B), pbs_resources_unicos8(7B), pbs_server_attributes(7B), pbs_resources_irix6(7B), pbs_resources_linux(7B).                                                                                                                                         | 9<br>9<br>9<br>9<br>9<br>9<br>9<br>9<br>9 |
| Example:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | 9                                         |
| <i>Int walltUse;</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | 9                                         |
| <i>// returns the amount of walltime used by</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | 9                                         |
| <i>// the job</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | 9                                         |
| <i>walltUse = JobIntResUseGet(job, "walltime");</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | 9                                         |
| <b>Size JobSizeResUseGet(Job job, String name)</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | 9                                         |
| Returns the value to resource specified in <b>name</b> used by the job (Job resources_used.name attribute). Call this function for resources with values that are of Size type. Sample resource names are: file, mem, pmem, workingset, pf, ppf, srfs_tmp, srfs_wrk, srfs_big, srfs_fast, sds, psds. For a description of these resource names, see pbs_resources_irix5(7B), pbs_resources_sp2(7B), pbs_resources_sunos4(7B), pbs_resources_unicos8(7B), pbs_server_attributes(7B), pbs_resources_irix6(7B), pbs_resources_linux(7B).                                                                                                                                                       | 9<br>9<br>9<br>9<br>9<br>9<br>9<br>9<br>9 |
| Example:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | 9                                         |
| <i>Size srfsUse;</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | 9                                         |
| <i>// returns the amount of srfs_fast used by</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | 9                                         |
| <i>// the job</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | 9                                         |
| <i>srfsUse = JobSizeResUseGet(job, "srfs_fast");</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | 9                                         |
| <b>String JobStringResUseGet(Job job, String name)</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | 9                                         |
| Returns the value to resource specified in <b>name</b> used by the job (Job resources_used.name attribute). Call this function for resources with values that are of String type. Sample resource names are: nodes, arch, neednodes. For a description of these resource names, see pbs_resources_irix5(7B), pbs_resources_sp2(7B), pbs_resources_sunos4(7B), pbs_resources_unicos8(7B), pbs_server_attributes(7B), pbs_resources_irix6(7B), pbs_resources_linux(7B).                                                                                                                                                                                                                       | 9<br>9<br>9<br>9<br>9<br>9<br>9<br>9      |
| (4) CNode-related functions                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | 9                                         |
| <b>Set CNode AllNodesGet(void)</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | 9                                         |
| Returns list of nodes that are managed by the server running on the local host. This could also include those nodes that were specified in the scheduler configuration file for which the scheduler writer wants the system to periodically check for information like state, property, and so on. See <b>pbs_sched_basl(8B)</b> for a discussion of configuration file format. CAUTION: This function must be called from inside <b>sched_main()</b> so that at every scheduling iteration, the most up to date <b>Set CNode</b> structure is returned. Do not call this from an assignment statement intended to initialize a global variable, as the statement will only be called once. | 9<br>9<br>9<br>9<br>9<br>9<br>9<br>9      |
| <b>CNode AllNodesLocalHostGet(void)</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | 9                                         |
| Returns the CNode object that represents the local host. This is a simple function to call for non-clustered systems where only 1 CNode exists. un-                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | 9<br>9                                    |



|                                                                                                                                                                                                                                                                                                                                                         |   |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| set value: NOCNODE                                                                                                                                                                                                                                                                                                                                      | 9 |
| CAUTION: This function must be called from inside <b>sched_main()</b> (or from within functions called by sched_main) so that at every scheduling iteration, the most up to date <b>CNode</b> structure is returned. Do not call this from an assignment statement intended to initialize a global variable, as the statement will only be called once. | 9 |
| <b>String CNodeNameGet(CNode node)</b>                                                                                                                                                                                                                                                                                                                  | 9 |
| Returns the unique (official) name of the node (i.e. ResMom hostname in a 1 mom/node model). This returns the same string that was specified in the configuration file. unset value: NULLSTR                                                                                                                                                            | 9 |
| <b>String CNodeOsGet(CNode node)</b>                                                                                                                                                                                                                                                                                                                    | 9 |
| Returns the os architecture of the node (i.e. "irix5", "sp2"). unset value: NULLSTR                                                                                                                                                                                                                                                                     | 9 |
| <b>Int CNodeStateGet( CNode node )</b>                                                                                                                                                                                                                                                                                                                  | 9 |
| Returns the node's state.                                                                                                                                                                                                                                                                                                                               | 9 |
| Return value:                                                                                                                                                                                                                                                                                                                                           | 9 |
| CNODE_OFFLINE, CNODE_DOWN, CNODE_FREE, CNODE_RE-                                                                                                                                                                                                                                                                                                        | 9 |
| SERVE, CNODE_INUSE_EXCLUSIVE, CN-                                                                                                                                                                                                                                                                                                                       | 9 |
| ODE_INUSE_SHARED, CNODE_UNKNOWN                                                                                                                                                                                                                                                                                                                         | 9 |
| <b>Int CNodeTypeGet( CNode node )</b>                                                                                                                                                                                                                                                                                                                   | 9 |
| Returns the node's type.                                                                                                                                                                                                                                                                                                                                | 9 |
| Return value:                                                                                                                                                                                                                                                                                                                                           | 9 |
| CNODE_TIMESHARED, CNODE_CLUSTER, CNODE_UN-                                                                                                                                                                                                                                                                                                              | 9 |
| KNOWN                                                                                                                                                                                                                                                                                                                                                   | 9 |
| <b>String CNodePropertiesGet(CNode node)</b>                                                                                                                                                                                                                                                                                                            | 9 |
| Returns the comma-separated list of other names the node is known by ( properties, other network name). For example, "babbage.nas.nasa.gov" maybe the node name, but it could also be known via "babbage1, babbage2". unset value: NULLSTR                                                                                                              | 9 |
| <b>String CNodeVendorGet(CNode node)</b>                                                                                                                                                                                                                                                                                                                | 9 |
| Returns the name of the vendor for the hardware of the machine (i.e. "sgi", "ibm"). unset value: NULLSTR                                                                                                                                                                                                                                                | 9 |
| <b>Int CNodeNumCpusGet(CNode node)</b>                                                                                                                                                                                                                                                                                                                  | 9 |
| Returns the number of processors attached to the node (or if it is a "job" node, then it's the # of cpus assigned to the job for this node). unset value: -1                                                                                                                                                                                            | 9 |
| <b>Int CNodeNumJobTasksGet(CNode node)</b>                                                                                                                                                                                                                                                                                                              | 9 |
| Returns the # of tasks (or processes) currently running on the "node". unset value: -1                                                                                                                                                                                                                                                                  | 9 |
| Example:                                                                                                                                                                                                                                                                                                                                                | 9 |
| <i>nodes = ServerNodesGet(s);</i>                                                                                                                                                                                                                                                                                                                       | 9 |
| <i>// Look at the system nodes</i>                                                                                                                                                                                                                                                                                                                      | 9 |
| <i>foreach(n in nodes) {</i>                                                                                                                                                                                                                                                                                                                            | 9 |
| <i>  print("SYSTEM NODES");</i>                                                                                                                                                                                                                                                                                                                         | 9 |
| <i>  print(n);</i>                                                                                                                                                                                                                                                                                                                                      | 9 |
| <i>  num_tasks = CNodeNumJobTasksGet(n);</i>                                                                                                                                                                                                                                                                                                            | 9 |
| <i>}</i>                                                                                                                                                                                                                                                                                                                                                | 9 |
| <i>The output will be:</i>                                                                                                                                                                                                                                                                                                                              | 9 |
| <i>SYSTEM NODES:</i>                                                                                                                                                                                                                                                                                                                                    | 9 |

|                                                                                                                                                                                                                                     |   |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| <i>Node Name = alabama</i>                                                                                                                                                                                                          | 9 |
| <i>NumJobTasks = 1</i>                                                                                                                                                                                                              | 9 |
| <i># of Processors = 1</i>                                                                                                                                                                                                          | 9 |
| ...                                                                                                                                                                                                                                 | 9 |
| <i>Node Name = aspasia</i>                                                                                                                                                                                                          | 9 |
| <i>NumJobTasks = 1</i>                                                                                                                                                                                                              | 9 |
| <i># of Processors = 1</i>                                                                                                                                                                                                          | 9 |
| ...                                                                                                                                                                                                                                 | 9 |
| <i>NOTE: "# of Processors" here means the<br/># of cpus (or virtual processors)<br/>residing on that node. The output<br/>of this call will always reflect<br/>the recent JobAction() of the<br/>scheduler (always up to date).</i> | 9 |
| <b>Size CNodeMemTotalGet( CNode node, String type )</b>                                                                                                                                                                             | 9 |
| Returns total memory of <b>type</b> for the node. <b>type</b> is an arbitrary string that the scheduler writer defines in the scheduler configuration file. unset value: -1b                                                        | 9 |
| Example:                                                                                                                                                                                                                            | 9 |
| <i>// get total physical memory</i>                                                                                                                                                                                                 | 9 |
| <i>CNodeMemTotalGet(node, "real")</i>                                                                                                                                                                                               | 9 |
| <i>// get total virtual memory</i>                                                                                                                                                                                                  | 9 |
| <i>CNodeMemTotalGet(node, "virtual")</i>                                                                                                                                                                                            | 9 |
| <b>Size CNodeMemAvailGet( CNode node, String type )</b>                                                                                                                                                                             | 9 |
| Returns available memory of <b>type</b> for the node. <b>type</b> is an arbitrary string that the scheduler writer defines in the scheduler configuration file. unset value: -1b                                                    | 9 |
| So sample calls will be:                                                                                                                                                                                                            | 9 |
| <i>// get available physical memory</i>                                                                                                                                                                                             | 9 |
| <i>CNodeMemAvailGet(node, "real")</i>                                                                                                                                                                                               | 9 |
| <i>// get available virtual memory</i>                                                                                                                                                                                              | 9 |
| <i>CNodeMemAvailGet(node, "virtual")</i>                                                                                                                                                                                            | 9 |
| <b>Int CNodeIdleTimeGet( CNode node )</b>                                                                                                                                                                                           | 9 |
| Returns number of seconds in which no keystroke or mouse movement has taken place on any terminal connected to the node. unset value: -1                                                                                            | 9 |
| <b>Float CNodeLoadAveGet( CNode node )</b>                                                                                                                                                                                          | 9 |
| Returns node's load average for all cpus. unset value: -1.0                                                                                                                                                                         | 9 |
| <b>Int CNodeCpuPercentIdleGet( CNode node )</b>                                                                                                                                                                                     | 9 |
| Returns the percent of idle time that all the processors of the node have experienced.                                                                                                                                              | 9 |
| <b>Int CNodeCpuPercentSysGet( CNode node )</b>                                                                                                                                                                                      | 9 |
| Returns the percent of time that all the processors of the node have spent running kernel code.                                                                                                                                     | 9 |
| <b>Int CNodeCpuPercentUserGet( CNode node )</b>                                                                                                                                                                                     | 9 |
| Returns the percent of time that all the processors of the node have spent running user code.                                                                                                                                       | 9 |
| <b>Int CNodeCpuPercentGuestGet( CNode node )</b>                                                                                                                                                                                    | 9 |
| Returns the percent of time that all the processors of the node have spent running a guest operating system.                                                                                                                        | 9 |
| <b>Int CNodeNetworkBwGet( CNode node, String type )</b>                                                                                                                                                                             | 9 |
| Returns the bandwidth of the node's network of <b>type</b> in bytes/second.                                                                                                                                                         | 9 |

|                                                                                                                                                                                                          |   |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| <b>type</b> is defined by the scheduler writer in the scheduler configuration file.                                                                                                                      | 9 |
| unset value: -1                                                                                                                                                                                          | 9 |
| Some sample calls are:                                                                                                                                                                                   | 9 |
| <i>CNodeNetworkBwGet( node, "hippi" );</i>                                                                                                                                                               | 9 |
| <i>CNodeNetworkBwGet( node, "fdi" );</i>                                                                                                                                                                 | 9 |
| <b>Size CNodeDiskSpaceTotalGet(CNode node, String name)</b>                                                                                                                                              | 9 |
| Returns the node's total space on disk identified by <b>name</b> where <b>name</b> is the device name arbitrarily defined by the scheduler writer in the scheduler configuration file. unset value: -1b  | 9 |
| Example:                                                                                                                                                                                                 | 9 |
| <i>CNodeDiskSpaceTotalGet( node, "/scratch2" );</i>                                                                                                                                                      | 9 |
| <b>Size CNodeDiskSpaceAvailGet(CNode node, String name)</b>                                                                                                                                              | 9 |
| Returns the node's available space on disk identified by <b>name</b> where <b>name</b> is arbitrarily defined by the scheduler writer in the scheduler configuration file. unset value: -1b              | 9 |
| Example:                                                                                                                                                                                                 | 9 |
| <i>CNodeDiskSpaceAvailGet( node, "/scratch1" );</i>                                                                                                                                                      | 9 |
| <b>Size CNodeDiskSpaceReservedGet(CNode node, String name)</b>                                                                                                                                           | 9 |
| Returns the node's reserved space on disk (user quota?) identified by <b>name</b> where <b>name</b> is arbitrarily defined by the scheduler writer in the scheduler configuration file. unset value: -1b | 9 |
| Example:                                                                                                                                                                                                 | 9 |
| <i>CNodeDiskSpaceReservedGet( node, "/scratch1" );</i>                                                                                                                                                   | 9 |
| <b>Int CNodeDiskInBwGet( CNode node, String name )</b>                                                                                                                                                   | 9 |
| Returns the write bandwidth (bytes/sec) of the node's disk identified by <b>name</b> . unset value: -1                                                                                                   | 9 |
| Example:                                                                                                                                                                                                 | 9 |
| <i>CNodeDiskInBwGet( node, "/fast" );</i>                                                                                                                                                                | 9 |
| <b>Int CNodeDiskOutBwGet( CNode node, String name )</b>                                                                                                                                                  | 9 |
| Returns read bandwidth (bytes/sec) of the node's disk identified by <b>name</b> . unset value: -1                                                                                                        | 9 |
| Example:                                                                                                                                                                                                 | 9 |
| <i>CNodeDiskOutBwGet( node, "/big" );</i>                                                                                                                                                                | 9 |
| <b>Size CNodeSwapSpaceTotalGet( CNode node, String name )</b>                                                                                                                                            | 9 |
| Returns the node's total space on swap identified by <b>name</b> where <b>name</b> is arbitrarily defined by the scheduler writer in the scheduler configuration file. unset value: -1b                  | 9 |
| Example:                                                                                                                                                                                                 | 9 |
| <i>CNodeSwapSpaceTotalGet( node, "primary" );</i>                                                                                                                                                        | 9 |
| <b>Size CNodeSwapSpaceAvailGet( CNode node, String name )</b>                                                                                                                                            | 9 |
| Returns node's available space on swap identified by <b>name</b> where <b>name</b> is the device name arbitrarily defined by the scheduler writer in the scheduler configuration file. unset value: -1b  | 9 |
| Example:                                                                                                                                                                                                 | 9 |
| <i>CNodeSwapSpaceAvailGet( node, "secondary" );</i>                                                                                                                                                      | 9 |
| <b>Int CNodeSwapInBwGet( CNode node, String name )</b>                                                                                                                                                   | 9 |
| Returns swapin rate of the node's swap device identified by <b>name</b> .                                                                                                                                | 9 |
| Example:                                                                                                                                                                                                 | 9 |
| <i>CNodeSwapInBwGet( node, "secondary" );</i>                                                                                                                                                            | 9 |
| <b>Int CNodeSwapOutBwGet( CNode node, String name )</b>                                                                                                                                                  | 9 |
| Returns the swapout rate of the node's swap device identified by <b>name</b> . unset value: -1                                                                                                           | 9 |

|                                                                                                                                                                                                                   |   |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| Example:                                                                                                                                                                                                          | 9 |
| <i>CNodeSwapOutBwGet(node, "primary");</i>                                                                                                                                                                        | 9 |
| <b>Size CNodeTapeSpaceTotalGet( CNode node, String name )</b>                                                                                                                                                     | 9 |
| Returns the node's total space on tape identified by <b>name</b> where <b>name</b> is arbitrarily defined by the scheduler writer in the scheduler configuration file. unset value: -1b                           | 9 |
| Example:                                                                                                                                                                                                          | 9 |
| <i>CNodeTapeSpaceTotalGet(node, "4mm");</i>                                                                                                                                                                       | 9 |
| <b>Size CNodeTapeSpaceAvailGet( CNode node, String name )</b>                                                                                                                                                     | 9 |
| Returns the node's available space on tape identified by <b>name</b> where <b>name</b> is arbitrarily defined by the scheduler writer in the scheduler configuration file. unset value: -1b                       | 9 |
| Example:                                                                                                                                                                                                          | 9 |
| <i>CNodeTapeSpaceAvailGet(node, "8mm");</i>                                                                                                                                                                       | 9 |
| <b>Int CNodeTapeInBwGet( CNode node, String name )</b>                                                                                                                                                            | 9 |
| Returns the write bandwidth (bytes/sec) of the node's tape identified by <b>name</b> . unset value: -1                                                                                                            | 9 |
| Example:                                                                                                                                                                                                          | 9 |
| <i>CNodeTapeInBwGet( node, "4mm" );</i>                                                                                                                                                                           | 9 |
| <b>Int CNodeTapeOutBwGet( CNode node, String name )</b>                                                                                                                                                           | 9 |
| Returns the read bandwidth (bytes/sec) of the node's tape identified by <b>name</b> . unset value: -1                                                                                                             | 9 |
| Example:                                                                                                                                                                                                          | 9 |
| <i>CNodeTapeOutBwGet( node, "8mm" );</i>                                                                                                                                                                          | 9 |
| <b>Size CNodeSrfsSpaceTotalGet( CNode node, String name )</b>                                                                                                                                                     | 9 |
| Returns the node's total space on srfs device identified by <b>name</b> where <b>name</b> is arbitrarily defined by the scheduler writer in the scheduler configuration file. unset value: -1b                    | 9 |
| Example:                                                                                                                                                                                                          | 9 |
| <i>CNodeSrfsSpaceTotalGet(node, "/fast");</i>                                                                                                                                                                     | 9 |
| <b>Size CNodeSrfsSpaceAvailGet( CNode node, String name )</b>                                                                                                                                                     | 9 |
| Returns the node's available space on srfs device identified by <b>name</b> where <b>name</b> is arbitrarily defined by the scheduler writer in some configuration file. unset value: -1b                         | 9 |
| Example:                                                                                                                                                                                                          | 9 |
| <i>CNodeSrfsSpaceAvailGet( node, "/big" );</i>                                                                                                                                                                    | 9 |
| <b>Size CNodeSrfsSpaceReservedGet(CNode node, String name)</b>                                                                                                                                                    | 9 |
| Returns the node's total amount of reserved space on srfs device identified by <b>name</b> where <b>name</b> is arbitrarily defined by the scheduler writer in the scheduler configuration file. unset value: -1b | 9 |
| Example:                                                                                                                                                                                                          | 9 |
| <i>CNodeSrfsSpaceReservedGet( node, "/fast" );</i>                                                                                                                                                                | 9 |
| <b>Int CNodeSrfsInBwGet( CNode node, String name )</b>                                                                                                                                                            | 9 |
| Returns the write bandwidth (bytes/sec) of the node's srfs device identified by <b>name</b> . unset value: -1                                                                                                     | 9 |
| Example:                                                                                                                                                                                                          | 9 |
| <i>CNodeSrfsInBwGet( node, "/fast" );</i>                                                                                                                                                                         | 9 |
| <b>Int CNodeSrfsOutBwGet( CNode node, String name )</b>                                                                                                                                                           | 9 |
| Returns the read bandwidth (bytes/sec) of the node's srfs device identified by <b>name</b> . unset value: -1                                                                                                      | 9 |
| Example:                                                                                                                                                                                                          | 9 |
| <i>CNodeSrfsOutBwGet( node, "/big" );</i>                                                                                                                                                                         | 9 |

|                                                                                                                               |   |
|-------------------------------------------------------------------------------------------------------------------------------|---|
| (5) Miscellaneous Functions                                                                                                   | 9 |
| <b>DateTime datetimeGet()</b>                                                                                                 | 9 |
| gets the current date/time.                                                                                                   | 9 |
| <b>Int datetimeToSecs(DateTime dt)</b>                                                                                        | 9 |
| returns the # of seconds since epoch (beginning of UNIX time - 00:00:00, January 1, 1970) for the given date/time <b>dt</b> . | 9 |
| <b>Int JobAction( Job job, Int action, String param )</b>                                                                     | 9 |
| Performs <b>action</b> on <b>job</b> with a <b>param</b> specified depending on the action.                                   | 9 |
| <b>action</b> can be: SYNCRUN, ASYNCRUN, DELETE, RERUN, HOLD, RELEASE, SIGNAL, MODIFYATTR, MODIFYRES where:                   | 9 |
| <i>Action</i> <i>Description</i>                                                                                              | 9 |
| =====                                                                                                                         | 9 |
| SYNCRUN <i>runs the job synchronously,</i>                                                                                    | 9 |
| <i>meaning the call to</i>                                                                                                    | 9 |
| <i>JobAction() will only</i>                                                                                                  | 9 |
| <i>return when the job has</i>                                                                                                | 9 |
| <i>started running or when</i>                                                                                                | 9 |
| <i>an error has been</i>                                                                                                      | 9 |
| <i>encountered.</i>                                                                                                           | 9 |
| <i>Param value:</i>                                                                                                           | 9 |
| <i>name of host(s) to run</i>                                                                                                 | 9 |
| <i>job under:</i>                                                                                                             | 9 |
| ASYNCRUN <i>runs the job asynchronously,</i>                                                                                  | 9 |
| <i>meaning the call to</i>                                                                                                    | 9 |
| <i>JobAction() will return</i>                                                                                                | 9 |
| <i>immediately as soon as</i>                                                                                                 | 9 |
| <i>the run request is</i>                                                                                                     | 9 |
| <i>validated by the PBS server,</i>                                                                                           | 9 |
| <i>and not necessarily when</i>                                                                                               | 9 |
| <i>the job has started</i>                                                                                                    | 9 |
| <i>execution.</i>                                                                                                             | 9 |
| <i>Param value:</i>                                                                                                           | 9 |
| <i>name of host(s) to run</i>                                                                                                 | 9 |
| <i>job under:</i>                                                                                                             | 9 |
| DELETE <i>deletes the job.</i>                                                                                                | 9 |
| <i>Param value:</i>                                                                                                           | 9 |
| <i>"deldelay=&lt;# of secs&gt;"</i>                                                                                           | 9 |
| <i>- delay # of seconds</i>                                                                                                   | 9 |
| <i>between the sending</i>                                                                                                    | 9 |
| <i>of SIGTERM and SIGKILL</i>                                                                                                 | 9 |
| <i>to the job before</i>                                                                                                      | 9 |
| <i>getting deleted.</i>                                                                                                       | 9 |
| RERUN <i>reruns the running job,</i>                                                                                          | 9 |
| <i>which involves terminating</i>                                                                                             | 9 |
| <i>the session leader of the</i>                                                                                              | 9 |
| <i>job and returning the job</i>                                                                                              | 9 |
| <i>to the queued state.</i>                                                                                                   | 9 |
| HOLD <i>places one or more holds</i>                                                                                          | 9 |
| <i>on the job.</i>                                                                                                            | 9 |
| <i>Param value:</i>                                                                                                           | 9 |

|                                                                                                                                                                  |                       |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------|
| "u", "o", "s", "uo", "os",<br>"uos"                                                                                                                              | 9<br>9                |
| - type of holds to place<br>on job: u(ser), o(ther),<br>s(ystem).                                                                                                | 9<br>9<br>9           |
| <i>RELEASE</i> removes or releases<br>holds placed on jobs.                                                                                                      | 9<br>9                |
| Param value:<br>"u", "o", "s", "uo", "os",<br>"uos"                                                                                                              | 9<br>9<br>9           |
| - type of holds to remove<br>from job: u(ser), o(ther),<br>s(ystem).                                                                                             | 9<br>9<br>9           |
| <i>SIGNAL</i> sends a signal to the<br>executing job.                                                                                                            | 9<br>9                |
| Param value:<br>"HUP", "SIGHUP",...                                                                                                                              | 9<br>9                |
| <i>MODIFYATTR</i> modifies the specified<br>attribute of the job to<br>the given value, when<br>the attrib_name is<br>!= "Resource_List" or<br>"resources_used". | 9<br>9<br>9<br>9<br>9 |
| Param value:<br>"attrib_name=value"                                                                                                                              | 9<br>9                |
| <i>MODIFYRES</i> modifies the job's<br>Resource_List<br>attribute given the<br>res_name and the<br>res_value:                                                    | 9<br>9<br>9<br>9      |
| Resource_List.res_name=<br>res_value                                                                                                                             | 9<br>9                |
| Param value:<br>"res_name=res_val"                                                                                                                               | 9<br>9                |
| <b>param</b> value depends on the action. Specify NULLSTR if no value for this<br>parameter is desired.                                                          | 9<br>9                |
| Return value: SUCCESS or FAIL.                                                                                                                                   | 9                     |
| NOTE: Any unrecognized <b>action</b> is ignored.                                                                                                                 | 9                     |
| Example:                                                                                                                                                         | 9                     |
| // run Job j synchronously<br>JobAction(j, SYNCRUN, NULLSTR);                                                                                                    | 9<br>9                |
| // run Job j asynchronously on host "db"<br>JobAction(j, ASYNCRUN, "db");                                                                                        | 9<br>9                |
| // delete Job j<br>JobAction(j, DELETE, NULLSTR);                                                                                                                | 9<br>9                |
| // delete Job j with a delay of 5 secs<br>// between the sending of SIGTERM and                                                                                  | 9<br>9                |

|                                                                                     |   |
|-------------------------------------------------------------------------------------|---|
| <i>// SIGKILL</i>                                                                   | 9 |
| <i>JobAction(j, DELETE, "deldelay=5");</i>                                          | 9 |
| <i>// rerun Job j</i>                                                               | 9 |
| <i>JobAction(j, RERUN, NULLSTR);</i>                                                | 9 |
| <i>// place a u(ser) hold on Job j</i>                                              | 9 |
| <i>JobAction(j, HOLD, "u");</i>                                                     | 9 |
| <i>// place an o(ther) hold on Job j</i>                                            | 9 |
| <i>JobAction(j, HOLD, "o");</i>                                                     | 9 |
| <i>// place a s(ystem) hold on Job j</i>                                            | 9 |
| <i>JobAction(j, HOLD, "s");</i>                                                     | 9 |
| <i>// place a default hold (u) on Job j</i>                                         | 9 |
| <i>JobAction(j, HOLD, NULLSTR);</i>                                                 | 9 |
| <i>// release u(ser) hold from Job j</i>                                            | 9 |
| <i>JobAction(j, RELEASE, "u");</i>                                                  | 9 |
| <i>// release o(ther) hold from Job j</i>                                           | 9 |
| <i>JobAction(j, RELEASE, "o");</i>                                                  | 9 |
| <i>// release s(ystem) hold from Job j</i>                                          | 9 |
| <i>JobAction(j, RELEASE, "s");</i>                                                  | 9 |
| <i>// release default hold (u) from Job j</i>                                       | 9 |
| <i>JobAction(j, RELEASE, NULLSTR);</i>                                              | 9 |
| <i>// send SIGHUP signal to Job j</i>                                               | 9 |
| <i>JobAction(j, SIGNAL, "SIGHUP");</i>                                              | 9 |
| <i>// update the comment attribute of Job</i>                                       | 9 |
| <i>// j to "a message".</i>                                                         | 9 |
| <i>// The param format is: attribute_name=new_value</i>                             | 9 |
| <i>// Consult PBS documentation for a list of job</i>                               | 9 |
| <i>// attribute names that can be specified.</i>                                    | 9 |
| <i>JobAction(j, MODIFYATTR, "comment=a message");</i>                               | 9 |
| <i>// update the Resource_List.cput attribute of Job</i>                            | 9 |
| <i>// j to 3600 seconds.</i>                                                        | 9 |
| <i>// The param format is: resource_name=new_value</i>                              | 9 |
| <i>// See pbs_resources* man page for a list of</i>                                 | 9 |
| <i>// resource_names that can be specified.</i>                                     | 9 |
| <i>JobAction(j, MODIFYRES, "cput=3600");</i>                                        | 9 |
| <b>QueJobFind(Que que, Fun Int func, Int cpr, Int value);</b>                       | 9 |
| <b>QueJobFind(Que que, Fun String func, Int cpr, String value);</b>                 | 9 |
| <b>QueJobFind(Que que, Fun DateTime func, Int cpr, DateTime value);</b>             | 9 |
| <b>QueJobFind(Que que, Fun Size func, Int cpr, Size value);</b>                     | 9 |
| where <b>cpr</b> is one of: OP_EQ, OP_NEQ, OP_LE, OP_LT, OP_GE, OP_GT.              | 9 |
| <b>func</b> is a function whose ONLY argument is of Job type. <b>Job</b> is the re- | 9 |
| turn type.                                                                          | 9 |

Description: Applies **func** to every job in **que**, and return the first job that satisfies the logical comparison: **func(job) cpr value**

Example: 9

```

Size JobVirtualMemAvailGet(Job job) 9
{ 9
    Size sz; 9

    sz = JobSizeResReqGet(job, "mem"); 9
    return(sz); 9
} 9
Int JobWallTimeReqGet(Job job) 9
{ 9
    Int wallt; 9

    wallt = JobIntResReqGet(job, "walltime"); 9
    return(wallt); 9
} 9
Int JobCpuTimeUsedGet(Job job) 9
{ 9
    Int cput; 9

    cput = JobIntResUseGet(job, "cput"); 9
    return(cput); 9
} 9
Que findQueByName(Set Que queues, String qname) 9
{ 9
    Que q; 9

    foreach(q in queues) { 9
        if( QueNameGet(q) EQ qname ) { 9
            return(q); 9
        } 9
    } 9
    return(NOQUE); 9
} 9
sched_main() 9
{ 9
    Server s; 9
    Que que; 9
    Set Que sq; 9

    // get local server 9
    s = AllServersLocalHostGet(); 9

    // get the queues of the Server s 9
    sq = ServerQueuesGet(s); 9

    // get the queue named "fast" from the 9
    // local server 9
    que = findQueByName( sq, "fast" ); 9

```



```

// Find the 1st job whose walltime requirement          9
// is == 300s:   9
QueJobFind(que, JobWallTimeReqGet, OP_EQ, 300);        9

// Find the 1st job whose email address to            9
// notify about job activity != "bayucan":            9
QueJobFind(que, JobEmailAddrGet, OP_NEQ,             9
           "bayucan");                                9

// Find the 1st job that was created after            9
// or on 3/3/1997:                                    9
QueJobFind(que, JobDateTimeCreatedGet, OP_GE,         9
           (3/3/1997));                               9

// Find the 1st job that was created after            9
// 3:3:44:   9
QueJobFind(que, JobDateTimeCreatedGet, OP_GT,         9
           (3:3:44));                                 9

// Find the 1st job that was created after            9
// 3:3:44 on 3/3/1997:                                9
QueJobFind(que, JobDateTimeCreatedGet, OP_GT,         9
           (3/3/1997@3:3:44));                       9

// Find the 1st job whose cpu time used < 1600s:     9
QueJobFind(que, JobCpuTimeUsedGet, OP_LT, 1600);     9

// Find the 1st job whose virtual memory             9
// requirement <= 300mb:                               9
QueJobFind(que, JobVirtualMemAvailGet, OP_LE,         9
           300mb);                                    9
}   9

Job QueJobFind( Que que, Fun Int func,   Int cpr)      9
Job QueJobFind( Que que, Fun String func, Int cpr)    9
Job QueJobFind( Que que, Fun DateTime func, Int cpr)  9
Job QueJobFind( Que que, Fun Size func,   Int cpr)    9

where cpr can be one of the following: OP_MAX, OP_MIN, func is a 9
function whose only argument is of Job type.           9

Description: Returns the Job with the max or min value found for 9
func(job) as it is applied to every job in que.    9

Example:   9
Int JobCpuTimeReqGet(Job job)                          9
{   9
    Int cput;   9

    cput = JobIntResReqGet(job, "cput");                9
    return(cput);                                       9
}   9
sched_main()   9
{   9

```

```

Que que; 9
Job job; 9

// Find the Job with the highest cpu time 9
// requirement: 9
job = QueJobFind(que, JobCpuTimeReqGet, OP_MAX); 9

// Find the Job with the minimum cpu time 9
// requirement: 9
job = QueJobFind(que, JobCpuTimeReqGet, OP_MIN); 9
} 9

Que QueFilter(Que que, Fun Int func, Int cpr, Int value) 9
Que QueFilter(Que que, Fun String func, Int cpr, String value) 9
Que QueFilter(Que que, Fun DateTime func, Int cpr, Date value) 9
Que QueFilter(Que que, Fun Size func, Int cpr, Size value) 9

where cpr can be one of the following: OP_EQ, OP_NEQ, OP_LE, OP_LT, 9
OP_GE, OP_GT, func is a function whose only argument is of Job type. 9

Description: Applies func to every job in que, and returns a new que con- 9
taining all jobs that satisfies the comparison condition: func(job) cpr 9
value 9

Example: 9
Int JobWallTimeReqGet(Job job) 9
{ 9
Int wallt; 9

wallt = JobIntResReqGet(job, "walltime"); 9
return(wallt); 9
} 9
sched_main() 9
{ 9
Que que; 9
Que newq; 9

// Returns a new que containing all jobs in "que" 9
// with a walltime requirement == 300s: 9
newq = QueFilter(que, JobWallTimeReqGet, OP_EQ, 300); 9

// Returns a new que containing all jobs in "que" 9
// with an email address != "bayucan": 9
newq = QueFilter(que, JobEmailAddrGet, OP_NEQ, "bayucan"); 9

// Returns a new que containing all jobs in "que" 9
// created after or on 3/3/1997: 9
newq = QueFilter(que, JobDateTimeCreatedGet, OP_GE, 9
(3/3/1997)); 9

// Returns a new que containing all jobs in "que" 9
// created after 3:3:44: 9
newq = QueFilter(que, JobDateTimeCreatedGet, OP_GT, 9
(3:3:44)); 9

```

```

// Returns a new que containing all jobs in "que"
// created after 3:3:44 on 3/3/1997:
newq = QueFilter(que, JobDateTimeCreatedGet, OP_GT,
                 (3|3|1997@3:3:44));

// NOTE: The original "que" is not modified
// whatsoever.
}

Int Sort(Set Job s, Fun Int key, Int order)
Int Sort(Set Job s, Fun String key, Int order)
Int Sort(Set Job s, Fun Float key, Int order)
Int Sort(Set Job s, Fun DateTime key, Int order)
Int Sort(Set Job s, Fun Size key, Int order)

where s the set of jobs to sort. key is the sorting key which is a function
whose only argument is of Job type, order is the sorting order: ASC, DE-
SC.

Description: sorts the elements of s, in either ASCending or DESCending
order of values that were returned by the key function, as applied to ev-
ery member of the set of jobs. The s object is modified with this call. This
returns SUCCESS or FAIL depending on outcome of the sort.

Examples:
Size JobMemReqGet(Job job)
{
    Size mem;

    mem = JobSizeResReqGet(job, "mem");
    return(mem);
}

sched_main()
{
    Server master;

    Set Job jobs;

    Int order;

    // get local server
    master = AllServersLocalHostGet();

    jobs = ServerJobsGet(master);
    Sort(jobs, JobPriorityGet, ASC);
    Sort(jobs, JobIdGet, DESC);
    order = ASC;
    Sort(jobs, JobDateTimeCreatedGet, order);
    order = DESC;
    Sort(jobs, JobMemReqGet, order);
}

Int Sort(Set Que s, Fun Int key, Int order)

```

**Int Sort(Set Que s, Fun String key, Int order)** 9  
**Int Sort(Set Que s, Fun Float key, Int order)** 9  
**Int Sort(Set Que s, Fun DateTime key, Int order)** 9  
**Int Sort(Set Que s, Fun Size key, Int order)** 9

where **s** the set of queues to sort. **key** is the sorting key which is a function whose only argument is of Que type, **order** is the sorting order: ASC, DESC. 9

Description: sorts the elements of **s**, in either ASCending or DESCending order of values that were returned by the **key** function, as applied to every member of the set of queues. The **s** object is modified with this call. This returns SUCCESS or FAIL depending on outcome of the sort. 9

Examples: 9

```
Size QueMemAvailGet(Que que) 9
{ 9
    Size mem; 9
    mem = QueSizeResAvailGet(que, "mem"); 9
    return(mem); 9
} 9
```

```
sched_main() 9
{ 9
    Server master; 9
    Set Que ques; 9
    Int order; 9
    // get local server 9
    master = AllServersLocalHostGet(); 9
    ques = ServerQueuesGet(master); 9
    Sort(ques, QuePriorityGet, ASC); 9
    Sort(ques, QueNameGet, ASC); 9
    order = DESC; 9
    Sort(ques, QueMemAvailGet, order); 9
} 9
```

**Int Sort(Set Server s, Fun Int key, Int order)** 9  
**Int Sort(Set Server s, Fun String key, Int order)** 9  
**Int Sort(Set Server s, Fun Float key, Int order)** 9  
**Int Sort(Set Server s, Fun DateTime key, Int order)** 9  
**Int Sort(Set Server s, Fun Size key, Int order)** 9

where **s** the set of servers to sort. **key** is the sorting key which is a function whose only argument is of Server type, **order** is the sorting order: ASC, DESC. 9

Description: sorts the elements of **s**, in either ASCending or DESCending order of values that were returned by the **key** function, as applied to every member of the set of servers. The **s** object is modified with this call. 9

|                                                                                                                                                                                                                                                                                                                   |   |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| This returns SUCCESS or FAIL depending on outcome of the sort.                                                                                                                                                                                                                                                    | 9 |
| Examples:                                                                                                                                                                                                                                                                                                         | 9 |
| <i>Size ServerMemAvailGet(Server serv)</i>                                                                                                                                                                                                                                                                        | 9 |
| {                                                                                                                                                                                                                                                                                                                 | 9 |
| <i>Size mem;</i>                                                                                                                                                                                                                                                                                                  | 9 |
| <br><i>mem = ServerSizeResAvailGet(serv, "mem");</i>                                                                                                                                                                                                                                                              | 9 |
| <i>return(mem);</i>                                                                                                                                                                                                                                                                                               | 9 |
| }                                                                                                                                                                                                                                                                                                                 | 9 |
| <br><i>sched_main()</i>                                                                                                                                                                                                                                                                                           | 9 |
| {                                                                                                                                                                                                                                                                                                                 | 9 |
| <i>Set Server sserver;</i>                                                                                                                                                                                                                                                                                        | 9 |
| <br><i>Int order;</i>                                                                                                                                                                                                                                                                                             | 9 |
| <br><i>Int ret;</i>                                                                                                                                                                                                                                                                                               | 9 |
| <br><i>sserver = AllServersGet();</i>                                                                                                                                                                                                                                                                             | 9 |
| <br><i>ret = Sort(sserver, ServerMaxRunJobsGet, ASC);</i>                                                                                                                                                                                                                                                         | 9 |
| <i>Sort(sserver, ServerInetAddrGet, ASC);</i>                                                                                                                                                                                                                                                                     | 9 |
| <br><i>order = DESC;</i>                                                                                                                                                                                                                                                                                          | 9 |
| <i>Sort(sserver, ServerMemAvailGet, order);</i>                                                                                                                                                                                                                                                                   | 9 |
| }                                                                                                                                                                                                                                                                                                                 | 9 |
| <b>Int Sort(Set CNode s, Fun Int key, Int order)</b>                                                                                                                                                                                                                                                              | 9 |
| <b>Int Sort(Set CNode s, Fun String key, Int order)</b>                                                                                                                                                                                                                                                           | 9 |
| <b>Int Sort(Set CNode s, Fun Float key, Int order)</b>                                                                                                                                                                                                                                                            | 9 |
| <b>Int Sort(Set CNode s, Fun DateTime key, Int order)</b>                                                                                                                                                                                                                                                         | 9 |
| <b>Int Sort(Set CNode s, Fun Size key, Int order)</b>                                                                                                                                                                                                                                                             | 9 |
| <br>where <b>s</b> the set of nodes to sort. <b>key</b> is the sorting key which is a function whose only argument is of CNode type, <b>order</b> is the sorting order: ASC, DESC.                                                                                                                                | 9 |
| <br>Description: sorts the elements of <b>s</b> , in either ASCending or DESCending order of values that were returned by the <b>key</b> function, as applied to every member of the set of nodes. The <b>s</b> object is modified with this call. This returns SUCCESS or FAIL depending on outcome of the sort. | 9 |
| Examples:                                                                                                                                                                                                                                                                                                         | 9 |
| <i>Size CNodeMyMemAvailGet(CNode cn)</i>                                                                                                                                                                                                                                                                          | 9 |
| {                                                                                                                                                                                                                                                                                                                 | 9 |
| <i>Size mem;</i>                                                                                                                                                                                                                                                                                                  | 9 |
| <br><i>mem = CNodeMemAvailGet(cn, "virtual");</i>                                                                                                                                                                                                                                                                 | 9 |
| <i>return(mem);</i>                                                                                                                                                                                                                                                                                               | 9 |
| }                                                                                                                                                                                                                                                                                                                 | 9 |
| <br><i>sched_main()</i>                                                                                                                                                                                                                                                                                           | 9 |
| {                                                                                                                                                                                                                                                                                                                 | 9 |

```

Set CNode scnode; 9
Int order; 9
scnode = AllNodesGet(); 9
Sort(scnode, CNodeIdleTimeGet, ASC); 9
Sort(scnode, CNodeNameGet, ASC); 9
order = DESC; 9
Sort(scnode, CNodeMyMemAvailGet, order); 9
} 9

```

**CNode..Get() FUNCTIONS** 9

The return values of the CNode..Get() functions discussed in the previous section 9  
are obtained by sending resource queries to the CNode's MOM at every scheduling 9  
iteration. For example, CNodeLoadAveGet(node) will return the value obtained 9  
from some <host resource> query (this could be the string "loadave") as sent to the 9  
node's MOM. The "<host resource> -> CNode..Get()" mappings are established in- 9  
ternally, but they can be modified or more mappings can be added via the scheduler 9  
configuration file. The config file is discussed in **pbs\_sched\_basl(8B)**. 9  
Mappings already established are given in the following: 9

For all architectures: 9

```

CNode..Get() actual call host resource 9
===== 9
CNodeOsGet(node) arch 9
CNodeLoadAveGet(node) loadave 9
CNodeIdleTimeGet(node) idletime 9

```

**SEE ALSO** 9

pbs\_sched\_basl(8B) pbs\_job\_attributes(7B), pbs\_queue\_attributes(7B), pbs\_re- 9  
sources\_iris5(7B), pbs\_resources\_sp2(7B), pbs\_resources\_sunos4(7B), pbs\_re- 9  
sources\_unicos8(7B), pbs\_server\_attributes(7B), and pbs\_server(8B), pbs\_re- 9  
sources\_iris6(7B), pbs\_resources\_linux(7B). 9

Figure 5-2 shows how the BASL data types Server, Set Server, Que, Set Que, Job, Set Job relate to one another. The items in italics are the predefined functions that must be called in a BASL program in order to instantiate the data types. The figure also illustrates the PBS framework: in a scheduling system, one or more server hosts exist, and each server manages one or more queues, and one or more jobs are spooled inside a queue. Jobs are run on one or more computational nodes (machines). A machine could act both as a server and a computation node.

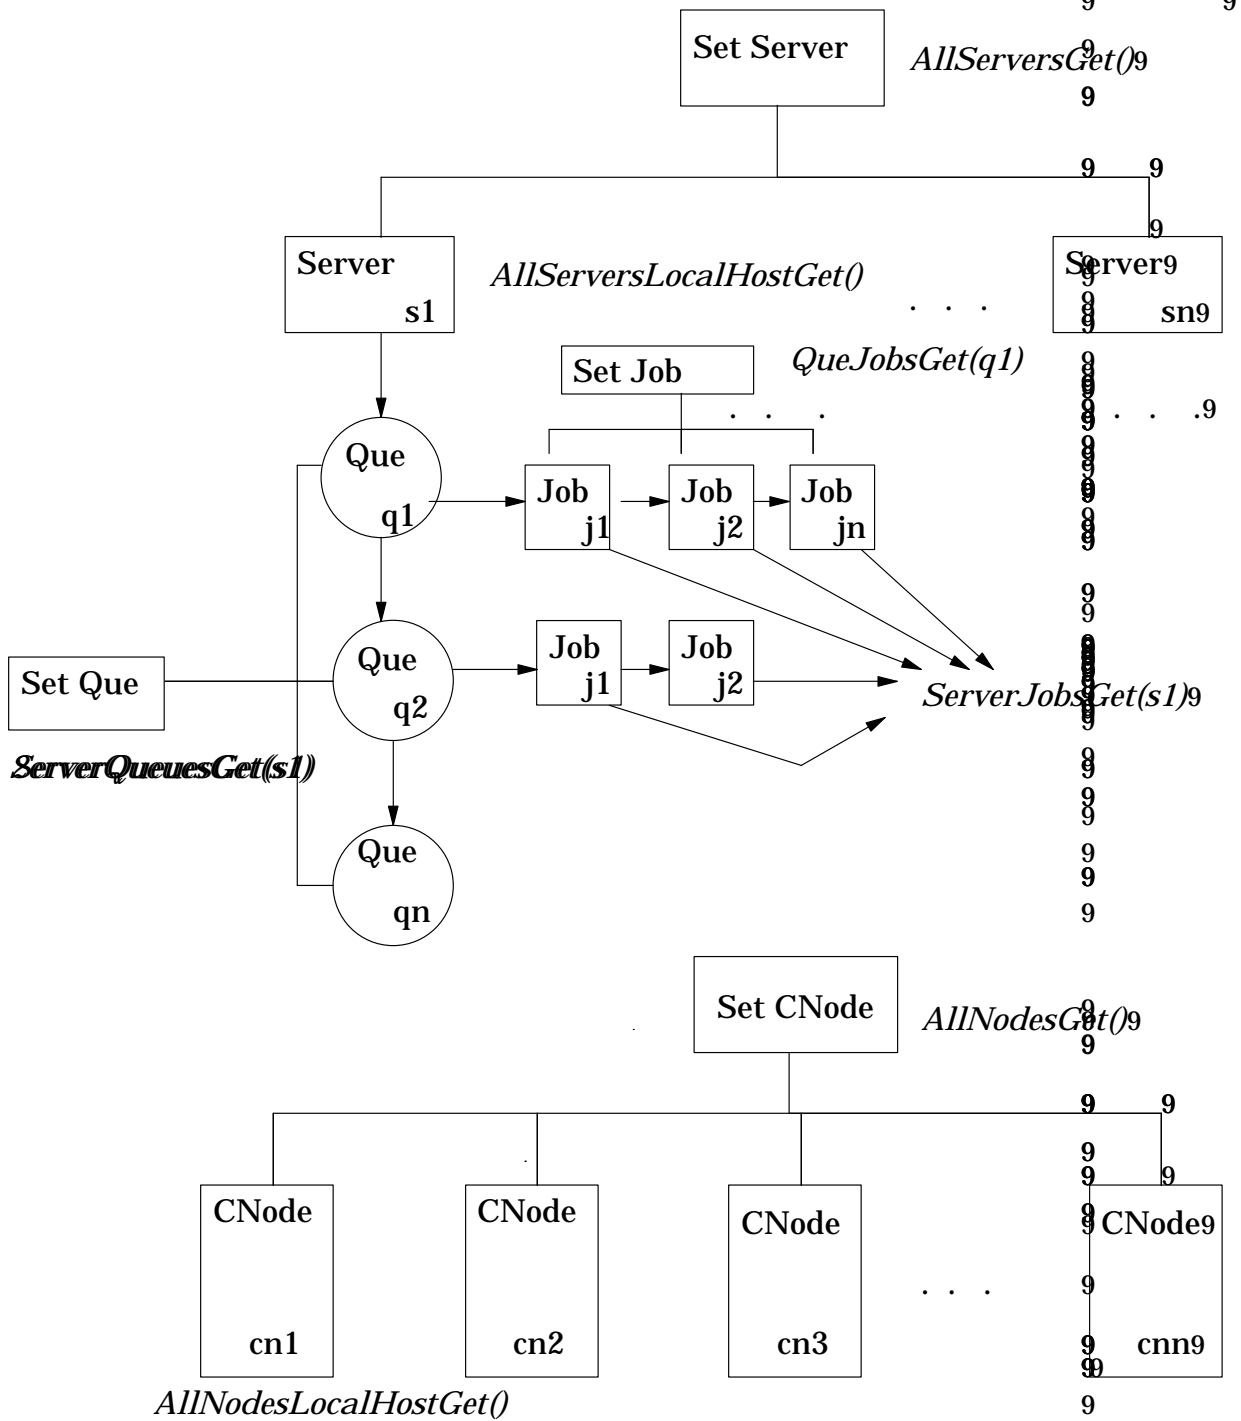


Figure 5-2: BASL Complex Data Types

### 5.2.15. Graphical User Interface: `xpbs` 9

**xpbs** is the graphical user interface to PBS user and operator commands. The current implementation will actually call the PBS commands like `qsub`, `qstat`, `qselect`, and so on. In addition, 2 new commands were created specifically for `xpbs`: `xpbs_scriptload`, and `xpbs_datadump`. These 2 binaries are compiled using the same libraries used by the standard PBS commands. 9

**xpbs** is written in Tcl/Tk. The way it works is that a main driver script/program called "xpbs" is what a user invokes, and this script calls other programs found in `XPBS_LIB` directory (as set in `Makefile`). 9

#### NAME 9

`xpbs` – GUI front end to PBS commands 9

#### SYNOPSIS 9

`xpbs [-admin]` 9

#### DESCRIPTION 9

The **xpbs** command provides a user-friendly point-and-click interface to PBS commands. Please see the sections below for a tour and tutorials. Also, within every dialog box, a **Help** button can be found for assistance. 9

#### OPTIONS 9

`-admin` A mode where additional buttons are made available for terminating PBS servers, starting/stopping/disabling/enabling queues, and running/rerunning jobs. 9

#### GETTING STARTED 9

Running **xpbs** will initialize the X resource database from various sources in the following order: 9

1. The **RESOURCE\_MANAGER** property on the root window (updated via `xrdb`) with settings usually defined in the `.Xdefaults` file 9
2. Preference settings defined by the system administrator in the global `xpbsrc` file 9
3. User's `~/xpbsrc` file - this file defines various X resources like fonts, colors, list of PBS hosts to query, criteria for listing queues and jobs, and various view states. See **PREFERENCES** section below for a list of resources that can be set. 9

#### RUNNING XPBS 9

To run **xpbs** as a regular, non-privileged user, type: 9

```
setenv DISPLAY <display_host>:0 9
xpbs 9
```

To run **xpbs** with the additional purpose of terminating PBS servers, stopping and starting queues, or running/rerunning jobs, then run: 9

```
xpbs -admin 9
```

NOTE: Be sure to appropriately set `~/rhosts` file if you're planning to submit jobs to some remote server, and expecting output files to be returned to the local host (where `xpbs` was run). Usually, adding the PBS hostname running the server to your `.rhosts` file locally, and adding the name of the local machine to the `.rhosts` file at remote host, should be sufficient. 9



Also, be sure that the PBS client commands are in the default PATH because **xpbs** will call these commands. 9

## THE XPBS DISPLAY 9

This section describes the main parts of the **xpbs** display. The main window is composed of 5 distinct areas (subwindows) arranged vertically (one on top of another) in the following order: 9

- 1) Menu 9
- 2) Hosts 9
- 3) Queues 9
- 4) Jobs 9
- 5) Info 9

**Menu.** The Menu area is composed of a row of command buttons that signal some action with a click of the left mouse button. The buttons are: 9

|               |                                                                                                                                                                                                                                                                                                                                                                         |   |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| Manual Update | to update the information on hosts, queues, and jobs.                                                                                                                                                                                                                                                                                                                   | 9 |
| Auto Update   | same as <i>Manual Update</i> except updating is done automatically every <some specified> number of minutes.                                                                                                                                                                                                                                                            | 9 |
| Track Job     | for periodically checking for returned output files of jobs.                                                                                                                                                                                                                                                                                                            | 9 |
| Preferences   | for setting certain parameters such as the list of server host(s) to query.                                                                                                                                                                                                                                                                                             | 9 |
| Help          | contains some help information.                                                                                                                                                                                                                                                                                                                                         | 9 |
| About         | tells of the author and who to send comments, bugs, suggestions to.                                                                                                                                                                                                                                                                                                     | 9 |
| Close         | for exiting <b>xpbs</b> plus saving the current setup information (if anything had changed) in the user's \$HOME/.xpbsrc file. Information saved include the selected host(s), queue(s), job(s), the different jobs listing criteria, the view states (i.e. minimized/maximized) of the Hosts, Queues, Jobs, and INFO regions, and anything in the Preferences section. | 9 |

**Hosts.** The Hosts area is composed of a leading horizontal HOSTS bar, a listbox, and a set of command buttons. The HOSTS bar contains a minimize/maximize button, identified by a dot or a rectangular image, for displaying or iconizing the Hosts region. The listbox displays information about favorite server host(s), and each entry is meant to be selected via a single left mouse button click, shift key + mouse button 1 click for contiguous selection, or cntrl key + mouse button 1 click for non-contiguous selection. The command buttons represent actions on selected host(s), and commonly found buttons are: 9

|           |                                                                                                                                                                |   |
|-----------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| detail    | for obtaining detailed information about selected server host(s). This functionality can also be achieved by double clicking on an entry in the Hosts listbox. | 9 |
| Submit    | for submitting a job to any of the queues managed by the selected host(s).                                                                                     | 9 |
| terminate | for terminating PBS servers on selected host(s). (-admin only)                                                                                                 | 9 |

The server hosts can be chosen by specifying in the ~/.xpbsrc file (or .Xdefaults) the resource: 9

```
*serverHosts: hostname1 hostname2 ... 9
```

Another way of specifying the host is to click on the Preferences button in the Menu region, and manipulate the server Hosts entry widget from the preferences dialog box. 9

**Queues.** The Queues area is composed of a leading horizontal QUEUES bar, a listbox, and a set of command buttons. The QUEUES bar lists the hosts that are consulted when listing queues; the bar also contains a minimize/maximize button for displaying or iconizing the 9

Queues region. The listbox displays information about queues managed by the server host(s) selected from the Hosts listbox; each listbox entry is meant to be selected (highlighted) via a single left mouse button click, shift key + mouse button 1 click for contiguous selection, or ctrl key + mouse button 1 click for non-contiguous selection. The command buttons represent actions for operating on selected queue(s), and commonly found buttons are:

- detail for obtaining detailed information about selected queue(s). This functionality can also be achieved by double clicking on a Queues listbox entry.
- stop for stopping the selected queue(s). (-admin only)
- start for starting the selected queue(s). (-admin only)
- disable for disabling the selected queue(s). (-admin only)
- enable for enabling the selected queue(s). (-admin only)

**Jobs.** The Jobs area is composed of a leading horizontal JOBS bar, a listbox, and a set of command buttons. The JOBS bar lists the queues that are consulted when listing jobs; the bar also contains a minimize/maximize button for displaying or iconizing the Jobs region. The listbox displays information about jobs that are found in the queue(s) selected from the Queues listbox; each listbox entry is meant to be selected (highlighted) via a single left mouse button click, shift key + mouse button 1 click for contiguous selection, or ctrl key + mouse button 1 click for non-contiguous selection. The region just above the Jobs listbox shows a collection of command buttons whose labels describe criteria used for filtering the Jobs listbox contents. The list of jobs can be selected according to the owner of jobs (Owners), job state (Job\_States), name of the job (Job\_Name), type of hold placed on the job (Hold\_Types), the account name associated with the job (Account\_Name), checkpoint attribute (Checkpoint), time the job is eligible for queueing/execution (Queue\_Time), resources requested by the job (Resources), priority attached to the job (Priority), and whether or not the job is rerunnable (Rerunnable). The selection criteria can be modified by clicking on any of the appropriate command buttons to bring up a selection box. The criteria command buttons are accompanied by a *Select Jobs* button, which when clicked, will update the contents of the Jobs listbox based on the new selection criteria. Please see **qselect(1B)** for more details on how the jobs are filtered.

Finally, to the right of the listbox, the Jobs region is accompanied by the following command buttons, for operating on selected job(s):

- detail for obtaining detailed information about selected job(s). This functionality can also be achieved by double clicking on a Jobs listbox entry.
- modify for modifying attributes of the selected job(s).
- delete for deleting the selected job(s).
- hold for placing some type of hold on selected job(s).
- release for releasing held job(s).
- signal for sending signals to selected job(s) that are running.
- msg for writing a message string into the output streams of the selected job(s).
- move for moving selected job(s) into some specified destination queue.
- order for exchanging order of two selected jobs in a queue.
- run for running selected job(s). (-admin only)
- rerun for requeueing selected job(s) that are running. (-admin only)

**Info.** The Info Area shows the progress of the commands' executed by **xpbs**. Look into this box for errors. The INFO bar also contains a minimize/maximize button for displaying or iconizing the Info region.

**WIDGETS USED IN XPBS**

Some of the widgets used in **xpbs** and how they are manipulated are described in the following:

1. **listbox** - can be multi-selectable (a number of entries can be selected/highlighted using a mouse click) or single-selectable (one entry can be highlighted at a time). For a multi-selectable listbox, the following operations are allowed:
  - a. single click with mouse button 1 to select/highlight an entry.
  - b. shift key + mouse button 1 to contiguously select more than one entry.
  - c. cntrl key + mouse button 1 to non-contiguously select more than one entry. NOTE: For systems running Tk < 4.0, the newly selected item is reshuffled to appear next to already selected items.
  - d. click the *Select All/Deselect All* button to select all entries or deselect all entries at once.
  - e. double clicking an entry usually activates some action that uses the selected entry as a parameter.
  
2. **scrollbar** - usually appears either vertically or horizontally and contains 5 distinct areas that are mouse clicked to achieve different effects:
 

|              |                                                                                                                                                                                                                                                                                          |
|--------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| top arrow    | Causes the view in the associated widget to shift up by one unit (i.e. the object appears to move down one unit in its window). If the button is held down the action will auto-repeat.                                                                                                  |
| top gap      | Causes the view in the associated window to shift up by one less than the number of units in the window (i.e. the portion of the object that used to appear at the very top of the window will now appear at the very bottom). If the button is held down the action will auto-repeat.   |
| slider       | Pressing button 1 in this area has no immediate effect except to cause the slider to appear sunken rather than raised. However, if the mouse is moved with the button down then the slider will be dragged, adjusting the view as the mouse is moved.                                    |
| bottom gap   | Causes the view in the associated window to shift down by one less than the number of units in the window (i.e. the portion of the object that used to appear at the very bottom of the window will now appear at the very top). If the button is held down the action will auto-repeat. |
| bottom arrow | Causes the view in the associated window to shift down by one unit (i.e. the object appears to move up one unit in its window). If the button is held down the action will auto-repeat.                                                                                                  |
  
3. **entry** - brought into focus with a click of the left mouse button. To manipulate this widget, simply type in the text value. Use of arrow keys, mouse selection of text for deletion or overwrite, copying and pasting with sole use of mouse buttons are permitted. This widget is usually accompanied by a scrollbar for horizontally scanning a long text entry string.
  
4. **matrix of entry boxes** - usually shown as several rows of entry widgets where a number of entries (called fields) can be found per row. The matrix is accompanied by up/down arrow buttons for paging through the rows of data, and each group of fields gets one scrollbar for horizontally scanning long entry strings. Moving from field to field can be done using the <Tab>, <Cntrl-f>, or <Cntrl-b> (move backwards) keys.

- 5. **spinbox** - a combination of an entry widget and a horizontal scrollbar. The entry widget will only accept values that fall within a defined list of valid values, and incrementing through the valid values is done by clicking on the up/down arrows. 9
- 6. **button** - a rectangular region appearing either raised or pressed that invokes an action when clicked with the left mouse button. When the button appears pressed, then hitting the <RETURN> key will automatically select the button. 9
- 7. **text** - an editor like widget. This widget is brought into focus with a click of the left mouse button. To manipulate this widget, simply type in the text. Use of arrow keys, backspace/delete key, mouse selection of text for deletion or overwrite, copying and pasting with sole use of mouse buttons are permitted. This widget is usually accompanied by a scrollbar for vertically scanning a long entry. 9

**SUBMITTING JOBS** 9

Submitting a PBS job requires only to manipulate the widgets found in the Submit window. The submit dialog box is composed of 4 distinct regions: 9

- 1) Job Script 9
- 2) OPTIONS 9
- 3) OTHER OPTIONS 9
- 4) Command Buttons 9

The Job Script file region is at the upper left, the OPTIONS region containing various widgets for setting job attributes is scattered all over the dialog box, the OTHER OPTIONS is located just below the Job Script file region, and Command Buttons region is at the bottom. 9

The job script region is composed of a header box, the text box, FILE entry box, and a couple of buttons labeled *load* and *save*. If you have a script file containing PBS options and executable lines, then type the name of the file on the FILE entry box, and then click on the *load* button. The various widgets in the Submit window will get loaded with values found in the script file. The script file text box will only be loaded with executable lines (non-PBS) found in the script. The job script header box has a *Prefix* entry box that can be modified to specify the PBS directive to look for when parsing a script file for PBS options. If you don't have a script file, you can start typing the executable lines of the job in the file text box. 9

To submit a job, perform the following steps: 9

- 1. Select a host from the HOSTS listbox in the main **xpbs** display. 9
- 2. Click on the *Submit* button located in the Menu bar. 9
- 3. Specify the script file containing the job execution lines and job property values, or simply type in the execution lines in the FILE textbox. 9
- 4. Start manipulating the various widgets in the Submit window. Particularly, pay close attention to the Destination listbox. This box lists all the queues found in the host that you selected. A special entry called "*@host*" refers to the default queue at host. Select appropriately the destination queue of the job. More options can be found by clicking the OTHER OPTIONS buttons. 9
- 5. At the bottom of the Submit window, click *confirm submit*. You can also click on *interactive* to run the job interactively. Running a job interactively will open an xterm window to your display host containing the session. 9

NOTE: The script FILE entry box is accompanied by a *save* button that you click to save the current widget values to the specified file in a form that can later be read by 9

- xpbs** or by the **qsub** command. 9
- MODIFYING ATTRIBUTES OF JOBS** 9
- Modifying a PBS job requires only to manipulate the widgets found in the Modify window. 9  
To modify a job or jobs, do the following steps: 9
1. Select one or more jobs from the JOBS listbox in the main **xpbs** display. 9
  2. Click on the *modify* button located to the right of the listbox. 9
  3. The Modify window is structured similarly to the Submit window. Simply manipulate the widgets to specify replacement or additional values of job attributes. 9
  4. Click on the *confirm modify* button located at the bottom of the dialog box. 9
- DELETING JOBS** 9
- Deleting a PBS job requires only to manipulate the widgets found in the Delete window. 9  
To delete a job or jobs, do the following steps: 9
1. Select one or more jobs from the JOBS listbox in the main **xpbs** display. 9
  2. Click on the *delete* button located to the right of the listbox. 9
  3. Manipulate the spinbox widget to set the kill delay signal interval. 9
  4. Click on the *delete* button located at the bottom of the dialog box. 9
- TRACKING RETURNED OUTPUT FILES** 9
- If you want to be informed of returned output files of current jobs, and be able to quickly see the contents of those files, then enable the "track job" feature as follows: 9
1. Submit all the jobs that you want monitored. 9
  2. Click on the *Track Job* button located in the Menu bar to bring up the Track Job dialog box. 9
  3. Specify the list of user names, whose jobs are to be monitored for returned output files, in the matrix located at the upper left of the dialog box. 9
  4. Manipulate the minutes spinbox, located just below the user names matrix, to specify the interval value when output files will be periodically checked. 9
  5. Specify the location of job output files (whether locally or remotely) by clicking on one of the radio buttons located at the upper right of the dialog box. Returned locally means the output files will be returned back to the host where **xpbs** was run. If the output files are returned to some remote host, then **xpbs** will execute an 9  
`RSH <remote_host> test -f <output_files>` 9  
to test the existence of the files. RSH is whatever you set the remote shell command 9  
to in the corresponding entry box. 9
- NOTE: Be sure the files are accessible from the host where **xpbs** was run (i.e. 9  
.rhosts appropriately set). 9
6. Click *start/reset tracking* button located at the bottom of the dialog box to: 9
    - cancel any previous tracking 9
    - build a new list of jobs to be monitored for returned output files based on currently 9  
queued jobs. 9
    - start periodic tracking. 9
  7. Click on *close window* button. 9
- When an output file for a job being monitored is found, then the *Track Job* button (the 9  
one that originally invoked the Track Job dialog box) will turn into a different color, and 9  
the *Jobs Found Completed* listbox, located in the Track Job dialog box, is then loaded 9

|                                                                                                                                                                                                                                                                                                                                                                                          |   |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| with the corresponding job id(s). Then double click on a job id to see the contents of the output file and the error file. Click <i>stop tracking</i> if you want to cancel tracking.                                                                                                                                                                                                    | 9 |
| <b>LEAVING XPBS</b>                                                                                                                                                                                                                                                                                                                                                                      | 9 |
| Click on the Close button located in the Menu bar to leave <b>xpbs</b> . If anything had changed, it will bring up a dialog box asking for a confirmation in regards to saving state information like the view states (minimize/maximize) of the HOSTS, QUEUES, JOBS, and INFO subwindows, and various criteria for listing queues and jobs. The information is saved in ~/.xpbsrc file. | 9 |
| <b>PREFERENCES</b>                                                                                                                                                                                                                                                                                                                                                                       | 9 |
| The resources that can be set in the X resources file, ~/.xpbsrc, are:                                                                                                                                                                                                                                                                                                                   | 9 |
| *serverHosts                                                                                                                                                                                                                                                                                                                                                                             | 9 |
| list of server hosts (space separated) to query by <b>xpbs</b> . A special keyword PBS_DEFAULT_SERVER can be used which will be used as a placeholder for the value obtained from *defServerFile.                                                                                                                                                                                        | 9 |
| *defServerFile                                                                                                                                                                                                                                                                                                                                                                           | 9 |
| the file containing the name of the default server host. The content of this will be substituted for the PBS_DEFAULT_SERVER keyword in *serverHosts value.                                                                                                                                                                                                                               | 9 |
| *timeoutSecs                                                                                                                                                                                                                                                                                                                                                                             | 9 |
| specify the number of seconds before timing out waiting for a connection to a PBS host.                                                                                                                                                                                                                                                                                                  | 9 |
| *xtermCmd                                                                                                                                                                                                                                                                                                                                                                                | 9 |
| the xterm command to run driving an interactive PBS session.                                                                                                                                                                                                                                                                                                                             | 9 |
| *labelFont                                                                                                                                                                                                                                                                                                                                                                               | 9 |
| font applied to text appearing in labels.                                                                                                                                                                                                                                                                                                                                                | 9 |
| *fixlabelFont                                                                                                                                                                                                                                                                                                                                                                            | 9 |
| font applied to text that label fixed-width widgets such as listbox labels. This must be a fixed-width font.                                                                                                                                                                                                                                                                             | 9 |
| *textFont                                                                                                                                                                                                                                                                                                                                                                                | 9 |
| font applied to a text widget. Keep this as fixed-width font.                                                                                                                                                                                                                                                                                                                            | 9 |
| *backgroundColor                                                                                                                                                                                                                                                                                                                                                                         | 9 |
| the color applied to background of frames, buttons, entries, scrollbar handles.                                                                                                                                                                                                                                                                                                          | 9 |
| *foregroundColor                                                                                                                                                                                                                                                                                                                                                                         | 9 |
| the color applied to text in any context (under selection, insertion, etc...).                                                                                                                                                                                                                                                                                                           | 9 |
| *activeColor                                                                                                                                                                                                                                                                                                                                                                             | 9 |
| the color applied to the background of a selection, a selected command button, or a selected scroll bar handle.                                                                                                                                                                                                                                                                          | 9 |
| *disabledColor                                                                                                                                                                                                                                                                                                                                                                           | 9 |
| color applied to a disabled widget.                                                                                                                                                                                                                                                                                                                                                      | 9 |
| *signalColor                                                                                                                                                                                                                                                                                                                                                                             | 9 |
| color applied to buttons that signal something to the user about a change of state. For example, the color of the <i>Track Job</i> button when returned output files are detected.                                                                                                                                                                                                       | 9 |
| *shadingColor                                                                                                                                                                                                                                                                                                                                                                            | 9 |
| a color shading applied to some of the frames to emphasize focus as well as decoration.                                                                                                                                                                                                                                                                                                  | 9 |
| *selectorColor                                                                                                                                                                                                                                                                                                                                                                           | 9 |
| the color applied to the selector box of a radiobutton or checkbutton.                                                                                                                                                                                                                                                                                                                   | 9 |
| *selectHosts                                                                                                                                                                                                                                                                                                                                                                             | 9 |
| list of hosts (space separated) to automatically select/highlight in the HOSTS listbox.                                                                                                                                                                                                                                                                                                  | 9 |
| *selectQueues                                                                                                                                                                                                                                                                                                                                                                            | 9 |
| list of queues (space separated) to automatically select/highlight in the QUEUES listbox.                                                                                                                                                                                                                                                                                                | 9 |

|                                                                                                                                                                                                                                                                                                     |   |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| *selectJobs                                                                                                                                                                                                                                                                                         | 9 |
| list of jobs (space separated) to automatically select/highlight in the JOBS listbox.                                                                                                                                                                                                               | 9 |
| *selectOwners                                                                                                                                                                                                                                                                                       | 9 |
| list of owners checked when limiting the jobs appearing on the Jobs listbox in the main <b>xpbs</b> window. Specify value as "Owners: <list_of_owners>". See -u option in <b>qselect(1B)</b> for format of <list_of_owners>.                                                                        | 9 |
| *selectStates                                                                                                                                                                                                                                                                                       | 9 |
| list of job states to look for (do not space separate) when limiting the jobs appearing on the Jobs listbox in the main <b>xpbs</b> window. Specify value as "Job_States: <states_string>". See -s option in <b>qselect(1B)</b> for format of <states_string>.                                      | 9 |
| *selectRes                                                                                                                                                                                                                                                                                          | 9 |
| list of resource amounts (space separated) to consult when limiting the jobs appearing on the Jobs listbox in the main <b>xpbs</b> window. Specify value as "Resources: <res_string>". See -l option in <b>qselect(1B)</b> for format of <res_string>.                                              | 9 |
| *selectExecTime                                                                                                                                                                                                                                                                                     | 9 |
| the Execution Time attribute to consult when limiting the list of jobs appearing on the Jobs listbox in the main <b>xpbs</b> window. Specify value as "Queue_Time: <exec_time>". See -a option in <b>qselect(1B)</b> for format of <exec_time>.                                                     | 9 |
| *selectAcctName                                                                                                                                                                                                                                                                                     | 9 |
| the name of the account that will be checked when limiting the jobs appearing on the Jobs listbox in the main <b>xpbs</b> window. Specify value as "Account_Name: <account_name>". See -A option in <b>qselect(1B)</b> for format of <account_name>.                                                | 9 |
| *selectCheckpoint                                                                                                                                                                                                                                                                                   | 9 |
| the checkpoint attribute relationship (including the logical operator) to consult when limiting the list of jobs appearing on the Jobs listbox in the main <b>xpbs</b> window. Specify value as "Checkpoint: <checkpoint_arg>". See -c option in <b>qselect(1B)</b> for format of <checkpoint_arg>. | 9 |
| *selectHold                                                                                                                                                                                                                                                                                         | 9 |
| the hold types string to look for in a job when limiting the jobs appearing on the Jobs listbox in the main <b>xpbs</b> window. Specify value as "Hold_Types: <hold_string>". See -h option in <b>qselect(1B)</b> for format of <hold_string>.                                                      | 9 |
| *selectPriority                                                                                                                                                                                                                                                                                     | 9 |
| the priority relationship (including the logical operator) to consult when limiting the list of jobs appearing on the Jobs listbox in the main <b>xpbs</b> window. Specify value as "Priority: <priority_value>". See -p option in <b>qselect(1B)</b> for format of <priority_value>.               | 9 |
| *selectRerun                                                                                                                                                                                                                                                                                        | 9 |
| the rerunnable attribute to consult when limiting the list of jobs appearing on the Jobs listbox in the main <b>xpbs</b> window. Specify value as "Rerunnable: <rerun_val>". See -r option in <b>qselect(1B)</b> for format of <rerun_val>.                                                         | 9 |
| *selectJobName                                                                                                                                                                                                                                                                                      | 9 |
| name of the job that will be checked when limiting the jobs appearing on the Jobs listbox in the main <b>xpbs</b> window. Specify value as "Job_Name: <jobname>". See -N option in <b>qselect(1B)</b> for format of <jobname>.                                                                      | 9 |
| *iconizeHostsView                                                                                                                                                                                                                                                                                   | 9 |
| a boolean value (true or false) indicating whether or not to iconize the HOSTS region.                                                                                                                                                                                                              | 9 |
| *iconizeQueuesView                                                                                                                                                                                                                                                                                  | 9 |
| a boolean value (true or false) indicating whether or not to iconize the QUEUES region.                                                                                                                                                                                                             | 9 |
| *iconizeJobsView                                                                                                                                                                                                                                                                                    | 9 |
| a boolean value (true or false) indicating whether or not to iconize the JOBS region.                                                                                                                                                                                                               | 9 |

|                                                                                                                                                                                                            |                                       |   |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------|---|
| <b>*iconizeInfoView</b>                                                                                                                                                                                    |                                       | 9 |
| a boolean value (true or false) indicating whether or not to iconize the INFO region.                                                                                                                      |                                       | 9 |
| <b>*jobResourceList</b>                                                                                                                                                                                    |                                       | 9 |
| a curly-braced list of resource names as according to architecture known to xpbs. The format is as follows:                                                                                                |                                       | 9 |
| { <arch-type1> resname1 resname2 ... resnameN }                                                                                                                                                            |                                       | 9 |
| { <arch-type2> resname1 resname2 ... resnameN }                                                                                                                                                            |                                       | 9 |
| ...                                                                                                                                                                                                        |                                       | 9 |
| { <arch-typeN> resname1 resname2 ... resnameN }                                                                                                                                                            |                                       | 9 |
| <b>XPBS AND PBS COMMANDS</b>                                                                                                                                                                               |                                       | 9 |
| <b>xpbs</b> calls PBS commands as follows:                                                                                                                                                                 |                                       | 9 |
| <b>Command Button</b>                                                                                                                                                                                      | <b>PBS Command</b>                    | 9 |
| detail (Hosts)                                                                                                                                                                                             | qstat -B -f <selected server_host(s)> | 9 |
| terminate                                                                                                                                                                                                  | qterm <selected server_host(s)>       | 9 |
| detail (Queues)                                                                                                                                                                                            | qstat -Q -f <selected queue(s)>       | 9 |
| stop                                                                                                                                                                                                       | qstop <selected queue(s)>             | 9 |
| start                                                                                                                                                                                                      | qstart <selected queue(s)>            | 9 |
| enable                                                                                                                                                                                                     | qenable <selected queue(s)>           | 9 |
| disable                                                                                                                                                                                                    | qdisable <selected queue(s)>          | 9 |
| detail (Jobs)                                                                                                                                                                                              | qstat -f <selected job(s)>            | 9 |
| modify                                                                                                                                                                                                     | qalter <selected job(s)>              | 9 |
| delete                                                                                                                                                                                                     | qdel <selected job(s)>                | 9 |
| hold                                                                                                                                                                                                       | qhold <selected job(s)>               | 9 |
| release                                                                                                                                                                                                    | qrls <selected job(s)>                | 9 |
| run                                                                                                                                                                                                        | qrun <selected job(s)>                | 9 |
| rerun                                                                                                                                                                                                      | qrerun <selected job(s)>              | 9 |
| signal                                                                                                                                                                                                     | qsig <selected job(s)>                | 9 |
| msg                                                                                                                                                                                                        | qmsg <selected job(s)>                | 9 |
| move                                                                                                                                                                                                       | qmove <selected job(s)>               | 9 |
| order                                                                                                                                                                                                      | qorder <selected job(s)>              | 9 |
| <b>EXIT STATUS</b>                                                                                                                                                                                         |                                       | 9 |
| Upon successful processing, the <b>xpbs</b> exit status will be a value of zero.                                                                                                                           |                                       | 9 |
| If the xpbs command fails, the command exits with a value greater than zero.                                                                                                                               |                                       | 9 |
| <b>SEE ALSO</b>                                                                                                                                                                                            |                                       | 9 |
| qalter(1B), qdel(1B), qhold(1B), qmove(1B), qmsg(1B), qrerun(1B), qrls(1B), qselect(1B), qsig(1B), qstat(1B), qorder(1B), qsub(1B), qdisable(8B), qenable(8B), qrun(8B), qstart(8B), qstop(8B), qterm(8B). |                                       | 9 |
| <b>5.2.16. Graphical User Interface: xpbsmon</b>                                                                                                                                                           |                                       | 9 |
| <b>xpbsmon</b> is the graphical user interface for displaying, monitoring the nodes/execution hosts under PBS.                                                                                             |                                       | 9 |
| <b>xpbsmon</b> is also written in Tcl/Tk. The way it works is that a main driver script/program                                                                                                            |                                       | 9 |



called "xpbsmon" is what a user invokes, and this script calls other programs found in XPB-SMON\_LIB directory (as set in Makefile). 9

## NAME 9

xpbsmon – GUI for displaying, monitoring the nodes/execution hosts under PBS 9

## SYNOPSIS 9

xpbsmon 9

## DESCRIPTION 9

The **xpbsmon** command provides a way to graphically display the various nodes that run 9 jobs. A node or execution host can be running a **pbs\_mom** daemon, or not running the dae- 9 mon. For the latter case, it could just be a nodename that appears in a nodes file that is man- 9 aged by a main **pbs\_server** running on another host. This utility also provides the ability to 9 monitor values of certain system resources by posting queries to the **pbs\_mom** of a node. 9 With this utility, you can see what job is running on what node, who owns the job, how many 9 nodes assigned to a job, status of each node (color-coded and the colors are user-modifiable), 9 how many nodes are available, free, down, reserved, offline, of unknown status, in use run- 9 ning multiple jobs or executing only 1 job. Please see the sections below for a tour and tutori- 9 als of xpbsmon. Also, within every dialog box, a **Help** button can be found for assistance. 9

## GETTING STARTED 9

Running xpbsmon will initialize the X resource database from various sources in the follow- 9 ing order: 9

1. The **RESOURCE\_MANAGER** property on the root window (updated via xrdb) with 9 settings usually defined in the .Xdefaults file 9
2. Preference settings defined by the system administrator in the global xpbsmonrc file 9
3. User's ~/.xpbsmonrc file - this file defines various X resources like fonts, colors, list of 9 colors to use to represent the various status of the nodes, list of PBS sites to query, list 9 of server hosts on each site, list of nodes/execution hosts on each server host, list of sys- 9 tem resource queries to send to the nodes' pbs\_mom, and various view states. See 9 PREFERENCES section below for a list of resources that can be set. 9

## RUNNING XPBSMON 9

**xpbsmon** can be run either as a regular user or superuser. If you run it with less privilege, 9 you may not be able to see all the information for a node. If it is executed as a regular user, 9 you should still be able to see what jobs are running on what nodes, possibly state and prop- 9 erties as these information are obtained by xpbsmon talking directly to the specified server. 9 If you want other system resource values, it may require special privilege since xpbsmon will 9 have to talk directly to the pbs\_mom of a node. In addition, the host where xpbsmon was run- 9 ning must also have been given explicit access permission by the mom (unless the GUI is 9 running on the same host where mom is running). This is done done by updating the \$lien- 9 thost and/or the \$restricted parameter on the mom's configuration file. 9

To run **xpbsmon**, type: 9

```
setenv DISPLAY <display_host>:0 9
xpbsmon 9
```

If you are running the GUI and only interested in jobs data, then be sure to set all the nodes' 9 type to NOMOM in the **Pref** dialog box. 9

## THE XPBSMON DISPLAY 9

This section describes the main parts of the **xpbsmon** display. The main window is composed of 3 distinct areas (subwindows) arranged vertically (one on top of another) in the following order:

- 1) Menu 9
- 2) Site Information 9
- 3) Info 9

**Menu.** The Menu area is composed of a row of command buttons that signal some action with a click of the left mouse button. The buttons are:

- Site.. displays a popup menu containing the list of PBS sites that have been added using the Sites Preferences window. Simply drag your mouse and release to the site name whose servers/nodes information you would like to see. 9
- Pref.. brings up various dialog boxes for specifying the list of sites, servers on each site, nodes that are known to a server, and the system resource queries to be sent to a node's pbs\_mom daemon. 9
- Auto Update.. brings up another window for specifying whether or not to do auto updates of nodes information, and also for specifying the interval number of minutes between updates. 9
- Help contains some help information. 9
- About tells who the author is and who to send comments, bugs, suggestions to. 9
- Close for exiting **xpbsmon** plus saving the current setup information (if anything had changed) in the user's \$HOME/.xpbsmonrc file. Information saved include the specified list of sites, servers on each site, nodes known to each server, and system resource queries to send to node's pbs\_mom. 9
- Minimize Button shows the iconized view of Site Information where nodes are represented as tiny boxes, where each box is colored according to status. In order to get more information about a node, you need to double click on the colored box. 9
- Maximize button shows the full view of Site Information where nodes are represented in bigger boxes, still colored depending on the status, and some information on it is displayed. 9

**Site Information.** Only one site at a time can be displayed. This area (shown as one huge box referred to as the site box) can be further sub-divided into 3 areas: the site name label at the top, server boxes in the middle, and the color status bar at the bottom. The site name label shows the name of the site as specified in the **Pref.** window. At the middle of the site box shows a row of big boxes housing smaller boxes. 9

The big box is an abstraction of a server host (called a server box), showing its server display label at the top of the box, a grid of smaller boxes representing the nodes that the server knows about (where jobs are run), and summary status for the nodes under the server. Status information will show counters for the number of nodes used, available, reserved, offline, or of unknown status and even # of cpus assigned. For a cleaner display, some counters with a value of zero are not displayed. The server boxes are placed in a grid, with a new row being started when either \*siteBoxMaxNumServerBoxesPerRow or \*siteBoxMaxWidth limit has been reached. 9

The smaller boxes represent the nodes/execution hosts where jobs are run (referred to as node boxes). Each node box shows the name at the top, and a sub-box (a smaller square) that is colored according to the status of the node that it represents, and if the view type is FULL, 9

it will display some node information according to the system resource queries specified on the **Pref.** window. Clicking on the sub-box will show a much bigger box (called the MIRROR view) with bigger fonts containing nodes information. Another view is called ICON and this shows a tiny box with a colored area. The node boxes are arranged in a grid, where a new row is created if either the `*serverBoxMaxNumNodeBoxesPerRow` or `*serverBoxMaxWidth` limit has been reached. ICON view of the node boxes will be constrained by the `*nodeBoxIconMaxHeight` and `*nodeBoxIconMaxWidth` pixel values; FULL view of the node boxes will be bounded by `*nodeBoxFullMaxWidth` and `*nodeBoxFullMaxHeight`; the mirror view of the node boxes has its size be `*nodeBoxMirrorMaxWidth`, and `*nodeBoxMirrorMaxHeight`.

Horizontal and vertical scrollbars for the site box, server box, and node box will be displayed as needed.

Finally, the color bar information shows a color chart displaying what the various colors mean in terms of node status. The color-to-status mapping can be modified by setting the X resources: `*nodeColorNOINFO`, `*nodeColorFREE`, `*nodeColorINUSEshared`, `*nodeColorINUSEexclusive`, `*nodeColorDOWN`, `*nodeColorRSVD`, `*nodeColorOFFL`.

**Info.** The Info Area shows the progress of some of the background actions performed by **xpbsmon**. Look into this box for errors.

#### WIDGETS USED IN XPBSMON

Some of the widgets used in **xpbsmon** and how they are manipulated are described in the following:

1. **listbox** - the ones found in this GUI are only single-selectable (one entry can be highlighted/selected at a time via a mouse click).
2. **scrollbar** - usually appears either vertically or horizontally and contains 5 distinct areas that are mouse clicked to achieve different effects:
 

|              |                                                                                                                                                                                                                                                                                          |
|--------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| top arrow    | Causes the view in the associated widget to shift up by one unit (i.e. the object appears to move down one unit in its window). If the button is held down the action will auto-repeat.                                                                                                  |
| top gap      | Causes the view in the associated window to shift up by one less than the number of units in the window (i.e. the portion of the object that used to appear at the very top of the window will now appear at the very bottom). If the button is held down the action will auto-repeat.   |
| slider       | Pressing button 1 in this area has no immediate effect except to cause the slider to appear sunken rather than raised. However, if the mouse is moved with the button down then the slider will be dragged, adjusting the view as the mouse is moved.                                    |
| bottom gap   | Causes the view in the associated window to shift down by one less than the number of units in the window (i.e. the portion of the object that used to appear at the very bottom of the window will now appear at the very top). If the button is held down the action will auto-repeat. |
| bottom arrow | Causes the view in the associated window to shift down by one unit (i.e. the object appears to move up one unit in its window). If the button is held down the action will auto-repeat.                                                                                                  |
3. **entry** - brought into focus with a click of the left mouse button. To manipulate this widget, simply type in the text value. Use of arrow keys, mouse selection of text for deletion or overwrite, copying and pasting with sole use of mouse buttons are permitted. This widget is usually accompanied by a scrollbar for horizontally scanning a long text entry

- string. 9
- 4. **box** - made up of 1 or more listboxes displayed adjacent to each other giving the effect of a "matrix". Each row from the listboxes makes up an element of the box. In order to add items to the box, you need to manipulate the accompanying entry widgets, one for each listbox, and then clicking the **add** button. Removing items from the box is done by selecting an element, and then clicking **delete**. 9
- 5. **spinbox** - a combination of an entry widget and a horizontal scrollbar. The entry widget will only accept values that fall within a defined list of valid values, and incrementing through the valid values is done by clicking on the up/down arrows. 9
- 6. **button** - a rectangular region appearing either raised or pressed that invokes an action when clicked with the left mouse button. When the button appears pressed, then hitting the <RETURN> key will automatically select the button. 9

**UPDATING PREFERENCES** 9

CASE 1: Time Sharing 9

Suppose you have a time-sharing environment where the front-end is called bower and you have 4 nodes: bower1, bower2, bower3, bower4. bower is the host that runs the server; jobs are submitted to host bower where it enqueues it for future execution. Also, a pbs\_mom daemon is running on each of the execution hosts. If the server bower also maintains a nodes list containing information like state, properties for the 4 nodes, then this will also be reported. Then to setup **xpbsmon**, do the following: 9

- 1. Click the **Pref.** button on the Menu section. 9
- 2. On the Sites Preference dialog, enter any arbitrary site name, for example "Local". Then click the **add** button. 9
- 3. On the Server\_Host entry box, enter "bower", and on the DisplayLabel entry box, put an arbitrary label (as it would appear on the header of the server box) like "Bower", and then click **add**. 9
- 4. Click the **nodes..** button that is accompanying the Servers box. This would bring up the Server Preference dialog. 9
- 5. Now add the entries "bower1", "bower2", "bower3", "bower4" specifying type MOM for each on the Nodes box. 9
- 6. If you need to monitor certain system resource parameters for each of the nodes, you need to specify query expressions containing resource queries to be sent to the individual PBS moms. For example, if you want to obtain memory usage, then select a node from the Nodes list, click on the **query..** button that accompanies the Nodes list, and this would bring up the Query Table dialog. Specify the following input: 9

Query\_Expr: (availmem/totmem) \* 100 9  
 Display\_Info: Memory Usage: 9  
 Display\_Type: SCALE 9

The above says to display the result of the "Query\_Expr" in a scale widget calibrated over 100. The queries "availmem" and "totmem" will be sent to the PBS mom, and the expression is evaluated upon receiving all results from the mom. If you want to display the result of another query, say "loadave", directly, then specify the following: 9

|                                                                                                                                                                                                                                                                                                                                                                                                                                              |   |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| Query_Expr: loadave                                                                                                                                                                                                                                                                                                                                                                                                                          | 9 |
| Display_Info: Load Average:                                                                                                                                                                                                                                                                                                                                                                                                                  | 9 |
| Display_Type: TEXT                                                                                                                                                                                                                                                                                                                                                                                                                           | 9 |
| NOTE: For a list of queries that can be sent to a pbs_mom, please click on the <b>Help</b> button on the Query table window.                                                                                                                                                                                                                                                                                                                 | 9 |
| CASE 2: Jobs Exclusive Environment                                                                                                                                                                                                                                                                                                                                                                                                           | 9 |
| Supposing you have a "space non-sharing" environment where the server maintains a list of nodes that it runs jobs on exclusively (one job at a time outstanding per node). Let's call this server b1. Simply update Preferences information as follows:                                                                                                                                                                                      | 9 |
| 1. Click the <b>Pref.</b> button on the Menu section.                                                                                                                                                                                                                                                                                                                                                                                        | 9 |
| 2. On the Sites Preference dialog, enter a site name, for example "B System". Then click the <b>add</b> button.                                                                                                                                                                                                                                                                                                                              | 9 |
| 3. On the Server_Host entry box, enter "b1", DisplayLabel entry box type "B1" (or whatever label that you would like to appear on the header of the server box), and then click <b>add</b> .                                                                                                                                                                                                                                                 | 9 |
| CASE 3: Hybrid Time Sharing/Space Sharing Environment                                                                                                                                                                                                                                                                                                                                                                                        | 9 |
| A cluster of heterogeneous machines, time-sharing or jobs exclusive, could easily be represented in <b>xpbsmon</b> by combining steps in CASE 1 and CASE 2.                                                                                                                                                                                                                                                                                  | 9 |
| <b>LEAVING XPBSMON</b>                                                                                                                                                                                                                                                                                                                                                                                                                       | 9 |
| Click on the <b>Close</b> button located in the Menu bar to leave <b>xpbsmon</b> . If anything had changed, it will bring up a dialog box asking for a confirmation in regards to saving preferences information about list of sites, their view types, list of servers on each site, the list of nodes known to each server, and the list of queries to be sent to the pbs_mom of each node. The information is saved in ~/.xpbsmonrc file. | 9 |
| <b>PREFERENCES</b>                                                                                                                                                                                                                                                                                                                                                                                                                           | 9 |
| The resources that can be set in the X resources file, ~/.xpbsmonrc, are described in the following:                                                                                                                                                                                                                                                                                                                                         | 9 |
| <b>Node Box Properties</b>                                                                                                                                                                                                                                                                                                                                                                                                                   | 9 |
| Resource names beginning with "*small" or "*node" apply to the properties of the node boxes. A node box is made of an outer frame where the node label sits on top, the canvas (smaller box) is on the middle, and possibly some horizontal/ vertical scrollbars.                                                                                                                                                                            | 9 |
| nodeColorNOINFO                                                                                                                                                                                                                                                                                                                                                                                                                              | 9 |
| color of node box when information for the node it represents could not be obtained.                                                                                                                                                                                                                                                                                                                                                         | 9 |
| *nodeColorFREE                                                                                                                                                                                                                                                                                                                                                                                                                               | 9 |
| color of canvas when node it represents is up.                                                                                                                                                                                                                                                                                                                                                                                               | 9 |
| *nodeColorINUSEshared                                                                                                                                                                                                                                                                                                                                                                                                                        | 9 |
| color when node it represents has more than 1 job running on it, or when node has been marked by the server that manages it as "job-sharing".                                                                                                                                                                                                                                                                                                | 9 |
| *nodeColorINUSEexclusive                                                                                                                                                                                                                                                                                                                                                                                                                     | 9 |
| list of colors to assign to a node box when host it represents is running only 1 job, or when node has been marked by the server that manages it as "time-sharing". xpbsmon will use this list to assign 1 distinct color per job unless all the colors have been exhausted, in which case, colors will start getting assigned more than once in a round-robin fashion.                                                                      | 9 |

|                                                                                           |   |
|-------------------------------------------------------------------------------------------|---|
| *nodeColorDOWN                                                                            | 9 |
| color when node it represents is down.                                                    | 9 |
| *nodeColorRSVD                                                                            | 9 |
| color when node it represents is reserved.                                                | 9 |
| *nodeColorOFFL                                                                            | 9 |
| color when node it represents is offline.                                                 | 9 |
| *smallForeground                                                                          | 9 |
| applies to the color of text inside the canvas.                                           | 9 |
| *smallBackground                                                                          | 9 |
| applies to the color of the frame.                                                        | 9 |
| *smallBorderWidth                                                                         | 9 |
| distance (in pixels) from other node boxes.                                               | 9 |
| *smallRelief                                                                              | 9 |
| how node box will visually appear (style).                                                | 9 |
| *smallScrollBorderWidth                                                                   | 9 |
| significant only in FULL mode, this is the distance of the horizontal/vertical scrollbars | 9 |
| from the canvas and lower edge of the frame.                                              | 9 |
| *smallScrollBackground                                                                    | 9 |
| background color of the scrollbars                                                        | 9 |
| *smallScrollRelief                                                                        | 9 |
| how scrollbars would visually appear (style).                                             | 9 |
| *smallCanvasBackground                                                                    | 9 |
| color of the canvas (later overridden depending on status of the node it represents)      | 9 |
| *smallCanvasBorderWidth                                                                   | 9 |
| distance of the canvas from the frame and possibly the scrollbars.                        | 9 |
| *smallCanvasRelief                                                                        | 9 |
| how the canvas is visually represented (style).                                           | 9 |
| *smallLabelBorderWidth                                                                    | 9 |
| the distance of the node label from the canvas and the topmost edge of the frame.         | 9 |
| *smallLabelBackground                                                                     | 9 |
| the background of the area of the node label that is not filled.                          | 9 |
| *smallLabelRelief                                                                         | 9 |
| how the label would appear visually (style).                                              | 9 |
| *smallLabelForeground                                                                     | 9 |
| the color of node label text.                                                             | 9 |
| *smallLabelFont                                                                           | 9 |
| the font to use for the node label text.                                                  | 9 |
| *smallLabelFontWidth                                                                      | 9 |
| font width (in pixels) of *smallLabelFont                                                 | 9 |
| *smallLabelFontHeight                                                                     | 9 |
| font height (in pixels) of *smallLabelFont                                                | 9 |
| *smallTextFont                                                                            | 9 |
| font to use for the text that appear inside a canvas.                                     | 9 |
| *smallTextFontWidth                                                                       | 9 |
| font width (in pixels) of *smallTextFont.                                                 | 9 |
| *smallTextFontHeight                                                                      | 9 |
| font height (in pixels) of *smallTextFont.                                                | 9 |

|                                                                                                                                                                                                                                                                                                           |   |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| *nodeColorTrough                                                                                                                                                                                                                                                                                          | 9 |
| color of trough part (the /100 portion) of a canvas scale item.                                                                                                                                                                                                                                           | 9 |
| *nodeColorSlider                                                                                                                                                                                                                                                                                          | 9 |
| color of slider part (value portion) of a canvas scale item.                                                                                                                                                                                                                                              | 9 |
| *nodeColorExtendedTrough                                                                                                                                                                                                                                                                                  | 9 |
| color of extended trough (over 100 portion when value exceeds max) of a canvas scale item.                                                                                                                                                                                                                | 9 |
| *nodeScaleFactor                                                                                                                                                                                                                                                                                          | 9 |
| tells how much bigger you want the scale item on the canvas to appear. (1 means to keep size as is)                                                                                                                                                                                                       | 9 |
| *nodeBoxFullMaxWidth                                                                                                                                                                                                                                                                                      | 9 |
| *nodeBoxFullMaxHeight                                                                                                                                                                                                                                                                                     | 9 |
| maximum width and height (in pixels) of a node box in FULL mode.                                                                                                                                                                                                                                          | 9 |
| *nodeBoxIconMaxWidth                                                                                                                                                                                                                                                                                      | 9 |
| *nodeBoxIconMaxHeight                                                                                                                                                                                                                                                                                     | 9 |
| maximum width and height (in pixels) of a node box in ICON mode.                                                                                                                                                                                                                                          | 9 |
| *nodeBoxMirrorMaxWidth                                                                                                                                                                                                                                                                                    | 9 |
| *nodeBoxMirrorMaxHeight                                                                                                                                                                                                                                                                                   | 9 |
| maximum width and height (in pixels) of a node box displayed on a separate window (after it has been clicked with the mouse to obtain a bigger view)                                                                                                                                                      | 9 |
| *nodeBoxMirrorScaleFactor                                                                                                                                                                                                                                                                                 | 9 |
| tells how much bigger you want the scale item on the canvas to appear while the node box is displayed on a separate window (1 means to keep size as is)                                                                                                                                                   | 9 |
| <b>Server Box Properties</b>                                                                                                                                                                                                                                                                              | 9 |
| Resource names beginning with "*medium" apply to the properties of the server boxes. A server box is made of an outer frame where the server display label sits on top, a canvas filled with node boxes is on the middle, possibly some horizontal/vertical scrollbars, and a status label at the bottom. | 9 |
| *mediumLabelForeground                                                                                                                                                                                                                                                                                    | 9 |
| color of text applied to the server display label and status label.                                                                                                                                                                                                                                       | 9 |
| *mediumLabelBackground                                                                                                                                                                                                                                                                                    | 9 |
| background color of the unfilled portions of the server display label and status label.                                                                                                                                                                                                                   | 9 |
| *mediumLabelBorderWidth                                                                                                                                                                                                                                                                                   | 9 |
| distance of the server display label and status label from other parts of the server box.                                                                                                                                                                                                                 | 9 |
| *mediumLabelRelief                                                                                                                                                                                                                                                                                        | 9 |
| how the server display label and status label appear visually (style).                                                                                                                                                                                                                                    | 9 |
| *mediumLabelFont"                                                                                                                                                                                                                                                                                         | 9 |
| font used for the text of the server display label and status label.                                                                                                                                                                                                                                      | 9 |
| *mediumLabelFontWidth                                                                                                                                                                                                                                                                                     | 9 |
| font width (in pixels) of *mediumLabelFont.                                                                                                                                                                                                                                                               | 9 |
| *mediumLabelFontHeight                                                                                                                                                                                                                                                                                    | 9 |
| font height (in pixels) of *mediumLabelFont.                                                                                                                                                                                                                                                              | 9 |
| *mediumCanvasBorderWidth                                                                                                                                                                                                                                                                                  | 9 |
| the distance of the server box's canvas from the label widgets.                                                                                                                                                                                                                                           | 9 |
| *mediumCanvasBackground                                                                                                                                                                                                                                                                                   | 9 |
| the background color of the canvas.                                                                                                                                                                                                                                                                       | 9 |

|                                                                                                  |   |
|--------------------------------------------------------------------------------------------------|---|
| *mediumCanvasRelief                                                                              | 9 |
| how the canvas appear visually (style).                                                          | 9 |
| *mediumScrollBorderWidth                                                                         | 9 |
| distance of the scrollbars from the other parts of the server box.                               | 9 |
| *mediumScrollBackground                                                                          | 9 |
| the background color of the scrollbars                                                           | 9 |
| *mediumScrollRelief                                                                              | 9 |
| how the scrollbars appear visually.                                                              | 9 |
| *mediumBackground                                                                                | 9 |
| the color of the server box frame.                                                               | 9 |
| *mediumBorderWidth                                                                               | 9 |
| the distance of the server box from other boxes.                                                 | 9 |
| *mediumRelief                                                                                    | 9 |
| how the server box appears visually (style).                                                     | 9 |
| *serverBoxMaxWidth                                                                               | 9 |
| *serverBoxMaxHeight                                                                              | 9 |
| maximum width and height (in pixels) of a server box.                                            | 9 |
| *serverBoxMaxNumNodeBoxesPerRow                                                                  | 9 |
| maximum # of node boxes to appear in a row within a canvas.                                      | 9 |
| <b>Miscellaneous Properties</b>                                                                  | 9 |
| Resource names beginning with "*big" apply to the properties of a site box, as well as to wid- 9 |   |
| gets found outside of the server box and node box. This includes the dialog boxes that appear 9  |   |
| when the menu buttons of the main window are manipulated. The site box is the one that ap- 9     |   |
| pears on the main region of xpbsmon. 9                                                           |   |
| *bigBackground                                                                                   | 9 |
| background color of the outer layer of the main window.                                          | 9 |
| *bigForeground                                                                                   | 9 |
| color applied to regular text that appear outside of the node box and server box.                | 9 |
| *bigBorderWidth                                                                                  | 9 |
| distance of the site box from the menu area and the color information area.                      | 9 |
| *bigRelief                                                                                       | 9 |
| how the site box is visually represented (style)                                                 | 9 |
| *bigActiveColor                                                                                  | 9 |
| the color applied to the background of a selection, a selected command button, or a se- 9        |   |
| lected scroll bar handle. 9                                                                      |   |
| *bigShadingColor                                                                                 | 9 |
| a color shading applied to some of the frames to emphasize focus as well as decoration. 9        |   |
| *bigSelectorColor                                                                                | 9 |
| the color applied to the selector box of a radiobutton or checkbutton. 9                         |   |
| *bigDisabledColor                                                                                | 9 |
| color applied to a disabled widget. 9                                                            |   |
| *bigLabelBackground                                                                              | 9 |
| color applied to the unfilled portions of label widgets. 9                                       |   |
| *bigLabelBorderWidth                                                                             | 9 |
| distance from other widgets of a label widget. 9                                                 |   |
| *bigLabelRelief                                                                                  | 9 |
| how label widgets appear visually (style) 9                                                      |   |



|                                                                                                                                         |   |
|-----------------------------------------------------------------------------------------------------------------------------------------|---|
| <b>*bigLabelFont</b>                                                                                                                    | 9 |
| font to use for labels.                                                                                                                 | 9 |
| <b>*bigLabelFontWidth</b>                                                                                                               | 9 |
| font width (in pixels) of <b>*bigLabelFont</b> .                                                                                        | 9 |
| <b>*bigLabelFontHeight</b>                                                                                                              | 9 |
| font height (in pixels) of <b>*bigLabelFont</b> .                                                                                       | 9 |
| <b>*bigLabelForeground</b>                                                                                                              | 9 |
| color applied to text that function as labels.                                                                                          | 9 |
| <b>*bigCanvasBackground</b>                                                                                                             | 9 |
| the color of the main region.                                                                                                           | 9 |
| <b>*bigCanvasRelief</b>                                                                                                                 | 9 |
| how the main region looks like visually (style)                                                                                         | 9 |
| <b>*bigCanvasBorderWidth:</b>                                                                                                           | 9 |
| distance of the main region from the menu and info regions.                                                                             | 9 |
| <b>*bigScrollBorderWidth</b>                                                                                                            | 9 |
| if the main region has a scrollbar, this is its distance from other widgets appearing on the region.                                    | 9 |
| <b>*bigScrollBackground</b>                                                                                                             | 9 |
| background color of the scrollbar appearing outside a server box and node box.                                                          | 9 |
| <b>*bigScrollRelief</b>                                                                                                                 | 9 |
| how the scrollbar that appears outside a server box and node box looks like visually (style)                                            | 9 |
| <b>*bigTextFontWidth</b>                                                                                                                | 9 |
| the font width (in pixels) of <b>*bigTextFont</b>                                                                                       | 9 |
| <b>*bigTextFontHeight</b>                                                                                                               | 9 |
| the font height (in pixels) of <b>*bigTextFont</b>                                                                                      | 9 |
| <b>*siteBoxMaxWidth</b>                                                                                                                 | 9 |
| maximum width (in pixels) of the site box.                                                                                              | 9 |
| <b>*siteBoxMaxHeight</b>                                                                                                                | 9 |
| maximum height (in pixels) of the site box.                                                                                             | 9 |
| <b>*siteBoxMaxNumServerBoxesPerRow</b>                                                                                                  | 9 |
| maximum number of server boxes to appear in a row inside the site box.                                                                  | 9 |
| <b>*autoUpdate</b>                                                                                                                      | 9 |
| if set to true, then information about nodes is periodically gathered.                                                                  | 9 |
| <b>*autoUpdateMins</b>                                                                                                                  | 9 |
| the # of minutes between polling for data regarding nodes when <b>*autoUpdate</b> is set.                                               | 9 |
| <b>*siteInView</b>                                                                                                                      | 9 |
| the name of the site that should be in view                                                                                             | 9 |
| <b>*rcSiteInfoDelimiterChar</b>                                                                                                         | 9 |
| the separator character for each input within a curly-bracketed line of input of <b>*siteInfo</b> .                                     | 9 |
| <b>*sitesInfo</b>                                                                                                                       | 9 |
| {<site1name><sep><site1-display-type><sep><server-name><sep><server-display-label><sep><nodename><sep><nodetype><sep><node-query-expr>} | 9 |
| ...                                                                                                                                     | 9 |
| {<site2name><sep><site2-display-type><sep><server-name><sep><server-display-label><sep><nodename><sep><nodetype><sep><node-query-expr>} | 9 |
| information about a site where <site1-display-type> can be either {FULL,ICON}, <node-                                                   | 9 |

type> can be {MOM, NOMOM}, and <node-query-expr> has the format: 9

{ {<expr>} {expr-label} <output-format>} 9

where <output-format> could be {TEXT, SCALE}. It's probably better to use the **Pref** dialog boxes in order to specify a value for this. 9

Example: 9

```
*rcSiteInfoDelimiterChar ; 9
*sitesInfo: {NAS;ICON;newton;Newton;newton3;NOMOM;} {Lang- 9
ley;FULL;db;DB;db.nas.nasa.gov;MOM;{{ ( availmem / totmem ) * 100} {Memory Us- 9
age;} SCALE} {{ ( loadave / ncpus ) * 100} {Cpu Usage;} SCALE} {ncpus {Number of 9
Cpus;} TEXT} {physmem {Physical Memory;} TEXT} {idletime {Idle Time (s);} TEXT} 9
{loadave {Load Avg;} TEXT}} {NAS;ICON;newton;Newton;newton4;NOMOM;} 9
{NAS;ICON;newton;Newton;newton1;NOMOM;} {NAS;ICON;newton;Newton;new- 9
ton2;NOMOM;} {NAS;ICON;b0101;DB;aspasia.nas.nasa.gov;MOM;{{ ( availmem / tot- 9
mem ) * 100} {Memory Usage;} SCALE} {{ ( loadave / ncpus ) * 100} {Cpu Usage;} 9
SCALE} {ncpus {Number of Cpus;} TEXT} {physmem {Physical Memory;} TEXT} {idle- 9
time {Idle Time (s);} TEXT} {loadave {Load Avg;} TEXT}} {NAS;ICON;newton;New- 9
ton;newton7;NOMOM;} 9
```

**EXIT STATUS** 9

Upon successful processing, the **xpbsmon** exit status will be a value of zero. 9

If the xpbsmon command fails, the command exits with a value greater than zero. 9

If xpbsmon is querying a host running a server with an incompatible version, you may see the following messages: 9

Internal error: pbsstatnode: End of File (15031) 9

The above message can be safely ignored. 9

**SEE ALSO** 9

pbs\_sched\_tcl(8B), pbs\_mom(8B). 9

[Page intentionally left blank.]

**6. Operator Commands**

A batch operator is a user of the system who is granted privilege to perform actions beyond those available to the normal user. Typical of those actions are starting and stopping of individual queues and the server itself, and modification or deleting the jobs of other users.

The operator will be able to use the following general user commands with increased privilege: qalter, qdel, qhold, qmove, qmsg, qrerun, qrls, qsig, and qstat. Additionally, the operator will have access to the new commands described in this section: qdisable, qenable, qrun, qstart, qstop, qterm, pbs\_mom, pbs\_sched, pbs\_server, and pbsnodes.

**6.1. Initiate Server Operation** 9**NAME** 9

**pbs\_server** – start a pbs batch server 9

**SYNOPSIS** 9

```
pbs_server [-a active] [-d config_path] [-e mask] [-p port] [-A acctfile] [-L logfile]
[-M mom_port] [-R momRPP_port] [-S scheduler_port] [-t type] [-g mom_globus_port]
[-G mom_globusRPP_port] 9
```

**DESCRIPTION** 9

The **pbs\_server** command starts the operation of a batch server on the local host. Typically, this command will be in a local boot file such as */etc/rc.local*. If the batch server is already in execution, **pbs\_server** will exit with an error. To insure that the **pbs\_server** command is not runnable by the general user community, the server will only execute if its real and effective uid is zero. 9

The server will record a diagnostic message in a log file for any error occurrence. The log files are maintained in the *server\_logs* directory below the home directory of the server. If the log file cannot be opened, the diagnostic message is written to the system console. 9

**OPTIONS** 9

**-a active** Specifies if scheduling is active or not. This sets the server attribute *scheduling*. If the option argument is "true" ("True", "t", "T", or "1"), the server is **active** and the PBS job scheduler will be called. If the argument is "false" ("False", "f", "F", or "0"), the server is **idle**, and the scheduler will not be called and no jobs will be run. If this option is not specified, the server will retain the prior value of the *scheduling* attribute. 9

**-d config\_path** Specifies the path of the directory which is home to the servers configuration files, *PBS\_HOME*. A host may have multiple servers. Each server must have a different configuration directory. The default configuration directory is given by the symbol *\$PBS\_SERVER\_HOME* which is typically */usr/spool/PBS*. 9

**-e mask** Specifies a log event mask to be used when logging. See "log\_events" in the *pbs\_server\_attributes(7B)* man page and in the ERS. 9

**-p port** Specifies the port number on which the server will listen for batch requests. If multiple servers are running on a single host, each must have its own unique port number. This option is for use in testing with multiple batch systems on a single host. 9

**-A acctfile** Specifies an absolute path name of the file to use as the accounting file. If not specified, the file is named for the current date in the *PBS\_HOME/server\_priv/accounting* directory. 9

**-L logfile** Specifies an absolute path name of the file to use as the log file. If not specified, the file is one named for the current date in the *PBS\_HOME/server\_logs* directory, see the **-d** option. 9

**-M mom\_port** Specifies the host name and/or port number on which the server should connect the job executor, MOM. The option argument, *mom\_conn*, is one of the forms: *host\_name*, *[:]port\_number*, or *host\_name:port\_number*. If *host\_name* not specified, the local host is assumed. If *port\_number* is not specified, the default port is assumed. See the **-M** option for *pbs\_mom(8)*. 9

|                               |                                                                                                                                                                                                                                                                                                                                                                                                                                                               |   |
|-------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| -R mom_RPPport                | Specifies the port number on which the server should query the up/down status of Mom. See the -R option for pbs_mom(8).                                                                                                                                                                                                                                                                                                                                       | 9 |
| -g mom_globus_port            | Specifies the host name and/or port number on which the server should connect the PBS Mom Globus daemon. The option argument, <i>mom_conn</i> , is one of the forms: <i>host_name</i> , <i>[:port_number]</i> , or <i>host_name:port_number</i> . If <i>host_name</i> not specified, the local host is assumed. If <i>port_number</i> is not specified, the default port is assumed.                                                                          | 9 |
| -G mom_globus_RPPport         | Specifies the port number on which the server should query the up/down status of PBS Mom Globus daemon.                                                                                                                                                                                                                                                                                                                                                       | 9 |
| -S scheduler_port             | Specifies the port number to which the server should connect when contacting the Scheduler. The option argument, <i>scheduler_conn</i> , is of the same syntax as under the -M option.                                                                                                                                                                                                                                                                        | 9 |
| -t type                       | Specifies the impact on jobs when the server restarts. <i>type</i> argument is:                                                                                                                                                                                                                                                                                                                                                                               | 9 |
| hot                           | All jobs in the Running state are retained in that state. Any job that was requeued into the Queued state from the Running state when the server last shut down will be run immediately, assuming the required resources are available. This returns the server to the same state as when it went down. After those jobs are restarted, then normal scheduling takes place for all remaining queued jobs. All other jobs are retained in their current state. | 9 |
|                               | If a job cannot be restarted immediately because of a missing resource, such as a node being down, the server will attempt to restart it periodically for upto 5 minutes. After that period, the server will revert to a normal state, as if warm started, and will no longer attempt to restart any remaining jobs which were running prior to the shutdown.                                                                                                 | 9 |
| warm                          | All jobs in the Running state are retained in that state. All other jobs are maintained in their current state. The job scheduler will typically make new selections for which jobs are placed into execution. Warm is the default if -t is not specified.                                                                                                                                                                                                    | 9 |
| cold                          | All jobs are purged. Positive confirmation is required before this direction is accepted.                                                                                                                                                                                                                                                                                                                                                                     | 9 |
| create                        | The server will discard any existing configuration files: <i>server</i> , <i>nodes</i> , <i>queues</i> and <i>jobs</i> , and initialize configuration files to the default values. The server is idled (scheduling set false).                                                                                                                                                                                                                                | 9 |
| <b>FILES</b>                  |                                                                                                                                                                                                                                                                                                                                                                                                                                                               | 9 |
| \$PBS_SERVER_HOME/server_priv | default directory for configuration files, typically <i>/usr/spool/pbs/server_priv</i>                                                                                                                                                                                                                                                                                                                                                                        | 9 |
| \$PBS_SERVER_HOME/server_logs | directory for log files recorded by the server.                                                                                                                                                                                                                                                                                                                                                                                                               | 9 |
| <b>Signal Handling</b>        |                                                                                                                                                                                                                                                                                                                                                                                                                                                               | 9 |
|                               | On receipt of the following signals, the server performs the defined action:                                                                                                                                                                                                                                                                                                                                                                                  | 9 |
| SIGHUP                        | The current server log and accounting log are closed and reopened. This allows for the prior log to be renamed and a new log started from the time of the signal.                                                                                                                                                                                                                                                                                             | 9 |

|                                                                                                                                                                                   |                                                                                                           |   |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------|---|
| SIGINT                                                                                                                                                                            | Causes an orderly "quick" shutdown of pbs_server, identical to "qterm".                                   | 9 |
| SIGTERM                                                                                                                                                                           | Causes an orderly "quick" shutdown of pbs_server, identical to "qterm".                                   | 9 |
| SIGSHUTDN                                                                                                                                                                         |                                                                                                           | 9 |
|                                                                                                                                                                                   | On systems (Unicos) where SIGSHUTDN is defined, it also causes an orderly "quick" shutdown of the server. | 9 |
| SIGPIPE, SIGUSR1, SIGUSR2                                                                                                                                                         |                                                                                                           | 9 |
|                                                                                                                                                                                   | These signals are ignored.                                                                                | 9 |
| All other signals                                                                                                                                                                 | have their default behavior installed.                                                                    | 9 |
| <b>EXIT STATUS</b>                                                                                                                                                                |                                                                                                           | 9 |
| If the server command fails to begin batch operation,                                                                                                                             | the server exits with a value greater than zero.                                                          | 9 |
| <b>SEE ALSO</b>                                                                                                                                                                   |                                                                                                           | 9 |
| qsub (1B), pbs_connect(3B), pbs_mom(8B), pbs_sched_basl(8B), pbs_sched_tcl(8B), pbsnodes(8B), qdisable(8B), qenable(8B), qmgr(8B), qrun(8B), qstart(8B), qstop(8B), and qterm(8B) |                                                                                                           | 9 |

**6.2. Initiate Execution Mini-server (MOM) Operation** 9

**NAME** 9

`pbs_mom` – start a pbs batch execution mini-server 9

**SYNOPSIS** 9

`pbs_mom [-C chkdirectory] [-c config] [-d directory] [-L logfile] [-M MOMport] [-R RPPport] [-p | -r] [-x] [-n nice_val]` 9

**DESCRIPTION** 9

The **pbs\_mom** command starts the operation of a batch **Machine Oriented Mini-server**, MOM, on the local host. Typically, this command will be in a local boot file such as */etc/rc.local*. To insure that the `pbs_mom` command is not runnable by the general user community, the server will only execute if its real and effective uid is zero. 9

One function of `pbs_mom` is to place jobs into execution as directed by the server, establish resource usage limits, monitor the job's usage, and notify the server when the job completes. If they exist, `pbs_mom` will execute a prologue script before executing a job and an epilogue script after executing the job. The next function of `pbs_mom` is to respond to resource monitor requests. This was done by a separate process in previous versions of PBS but has now been combined into one process. The resource monitor function is provided mainly for the PBS scheduler. It provides information about the status of running jobs, memory available etc. The next function of `pbs_mom` is to respond to task manager requests. This involves communicating with running tasks over a tcp socket as well as communicating with other MOMs within a job (aka a "sisterhood"). 9

`Pbs_mom` will record a diagnostic message in a log file for any error occurrence. The log files are maintained in the *mom\_logs* directory below the home directory of the server. If the log file cannot be opened, the diagnostic message is written to the system console. 9

If `pbs_mom` was built under the *irix6\_cpuset* architecture, then it will fork into 2 (at most 3) `pbs_moms`: a main mom that services job; a collector mom that gathers process information for every process that it tracks; and a hammer mom which is in charge of killing off any stray processes or unauthorized processes running on the system. The hammer mom is turned off by default. 9

**OPTIONS** 9

`-C chkdirectory` Specifies the path of the directory used to hold checkpoint files. [Currently this is only valid on Cray systems.] The default directory is `PBS_HOME/spool/checkpoint`, see the `-d` option. The directory specified with the `-C` option must be owned by root and accessible (`rwX`) only by root to protect the security of the checkpoint files. 9

`-c config` Specify a alternative configuration file, see description below. If this is a relative file name it will be relative to `PBS_HOME/mom_priv`, see the `-d` option. If the specified file cannot be opened, `pbs_mom` will abort. If the `-c` option is not supplied, `pbs_mom` will attempt to open the default configuration file "config" in `PBS_HOME/mom_priv`. If this file is not present, `pbs_mom` will log the fact and continue. 9

`-d directory` Specifies the path of the directory which is the home of the servers working files, `PBS_HOME`. This option is typically used along with `-M` when debugging MOM. The default directory is given by `$PBS_SERVER_HOME` which is typically `/usr/spool/PBS`. 9

`-L logfile` Specify an absolute path name for use as the log file. If not specified, MOM will open a file named for the current date in the `PBS_HOME/mom_logs` directory, see the `-d` option. 9



|             |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |                                      |
|-------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------|
| -M port     | Specifies the port number on which the mini-server (MOM) will listen for batch requests.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | 9<br>9                               |
| -R port     | Specifies the port number on which the mini-server (MOM) will listen for resource monitor requests, task manager requests and inter-MOM messages. Both a UDP and a TCP port of this number will be used.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | 9<br>9<br>9                          |
| -p          | Specifies the impact on jobs which were in execution when the mini-server shut down. On any restart of MOM, the new mini-server will not be the parent of any running jobs, MOM has lost control of her offspring (not a new situation for a mother). With the -p option, Mom will allow the jobs to continue to run and monitor them indirectly via polling. The -p option is mutually exclusive with the -r option.                                                                                                                                                                                                                                                                                                                                 | 9<br>9<br>9<br>9<br>9<br>9           |
| -r          | Specifies the impact on jobs which were in execution when the mini-server shut down. With the -r option, MOM will kill any processes belonging to jobs, mark the jobs as terminated, and notify the batch server which owns the job. The -r option is mutual exclusive with the -p option.<br><br>Normally the mini-server is started from the system boot file without the -p or the -r option. The mini-server will make no attempt to signal the former session of any job which may have been running when the mini-server terminated. It is assumed that on reboot, all processes have been killed.<br><br>If the -r option is used following a reboot, process ids (pids) may be reused and MOM may kill a process that is not a batch session. | 9<br>9<br>9<br>9<br>9<br>9<br>9<br>9 |
| -a alarm    | Used to specify the alarm timeout in seconds for computing a resource. Every time a resource request is processed, an alarm is set for the given amount of time. If the request has not completed before the given time, an alarm signal is generated. The default is 5 seconds.                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | 9<br>9<br>9<br>9                     |
| -x          | Disables the check for privileged port resource monitor connections. This is used mainly for testing since the privileged port is the only mechanism used to prevent any ordinary user from connecting.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | 9<br>9<br>9                          |
| -n nice_val | Specifies the priority value of the daemon when it executes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | 9                                    |

### CONFIGURATION FILE 9

The configuration file may be specified on the command line at program start with the -c flag. The use of this file is to provide several types of run time information to pbs\_mom: static resource names and values, external resources provided by a program to be run on request via a shell escape, and values to pass to internal set up functions at initialization (and re-initialization). 9  
9  
9  
9  
9

Each item type is on a single line with the component parts separated by white space. If the line starts with a hash mark (pound sign, #), the line is considered to be a comment and is skipped. 9  
9

#### Shell Commands 9

If the first character of the value is an exclamation mark (!), the entire rest of the line is saved to be executed through the services of the **system(3)** standard library routine. 9

The shell escape provides a means for the resource monitor to yield arbitrary information to the scheduler. Parameter substitution is done such that the value of any qualifier sent with the query, as explained below, replaces a token with a percent sign (%) followed by the name of the qualifier. For example, here is a configuration file line which gives a resource name of "escape": 9  
9  
9  
9

```
escape      !echo %xxx %yyy 9
```

If a query for "escape" is sent with no qualifiers, the command executed would be "echo %xxx %yyy". If one qualifier is sent, "escape[xxx=hi there]", the command executed 9  
9

would be "echo hi there %yyy". If two qualifiers are sent, "escape[xxx=hi][yyy=there]", the command executed would be "echo hi there". If a qualifier is sent with no matching token in the command line, "escape[zzz=snafu]", an error is reported.

**Initialization Value**

An initialization value directive has a name which starts with a dollar sign (\$) and must be known to MOM via an internal table. The entries in this table now are:

**clienthost**

which causes a host name to be added to the list of hosts which will be allowed to connect to MOM as long as they are using a privileged port. For example, here are two configuration file lines which will allow the hosts "fred" and "wilma" to connect:

```
$clienthost      fred
$clienthost      wilma
```

Two host name are always allowed to connection to pbs\_mom, "localhost" and the name returned to pbs\_mom by the system call gethostname(). These names need not be specified in the configuration file. The hosts listed as "clienthosts" comprise a "sisterhood" of machines. Any one of the sisterhood will accept connections from a server from within the sisterhood. They will also accept Resource Monitor (RM) requests and Internal MOM (IM) messages from within the sisterhood. For a sisterhood to be able to communicate IM messages to each other, they must all share the same RM port.

**restricted**

which causes a host name to be added to the list of hosts which will be allowed to connect to MOM without needing to use a privileged port. These names allow for wildcard matching. For example, here is a configuration file line which will allow queries from any host from the domain "ibm.com".

```
$restricted      *.ibm.com
```

The restriction which applies to these connections is that only internal queries may be made. No resources from a config file will be found. This is to prevent any shell commands from being run by a non-root process.

**logevent**

which sets the mask that determines which event types are logged by pbs\_mom. For example:

```
$logevent 0x1fff
$logevent 255
```

The first example would set the log event mask to 0x1ff (511) which enables logging of all events including debug events. The second example would set the mask to 0x0ff (255) which enables all events except debug events.

**cputmult**

which sets a factor used to adjust cpu time used by a job. This is provided to allow adjustment of time charged and limits enforced where the job might run on systems with different cpu performance. If Mom's system is faster than the reference system, set cputmult to a decimal value greater than 1.0. If Mom's system is slower, set cputmult to a value between 1.0 and 0.0. For example:

```
$cputmult 1.5
$cputmult 0.75
```

**wallmult**

which sets a factor used to adjust wall time usage by to job to a common reference system. The factor is used for walltime calculations and limits the same as cputmult is used for cpu time.

|                                                                                                                                                                                                                                                                                             |   |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| Static Resources                                                                                                                                                                                                                                                                            | 9 |
| For static resource names and values, the configuration file contains a list of resource names/values pairs, one pair per line and separated by white space. An Example of static resource names and values could be the number of tape drives of different types and could be specified by | 9 |
| tape3480           4                                                                                                                                                                                                                                                                        | 9 |
| tape3420           2                                                                                                                                                                                                                                                                        | 9 |
| tapedat            1                                                                                                                                                                                                                                                                        | 9 |
| tape8mm            1                                                                                                                                                                                                                                                                        | 9 |
| If pbs_mom was built under irix6_cpuset, then the following additional directives are available:                                                                                                                                                                                            | 9 |
| {enforce   use   restricted}                                                                                                                                                                                                                                                                | 9 |
| This signals to mom that the following limits should be enforced/enabled for a job:                                                                                                                                                                                                         | 9 |
| enforce mem                                                                                                                                                                                                                                                                                 | 9 |
| enforce !mem                                                                                                                                                                                                                                                                                | 9 |
| Enforce or don't enforce job "mem" request. This is enforced by default.                                                                                                                                                                                                                    | 9 |
| enforce pvmem                                                                                                                                                                                                                                                                               | 9 |
| enforce !pvmem                                                                                                                                                                                                                                                                              | 9 |
| Enforce or don't enforce job "pvmem" request. This is enforced by default.                                                                                                                                                                                                                  | 9 |
| enforce vmem                                                                                                                                                                                                                                                                                | 9 |
| enforce !vmem                                                                                                                                                                                                                                                                               | 9 |
| Enforce or don't enforce job "vmem" request. This is enforced by default.                                                                                                                                                                                                                   | 9 |
| enforce walltime                                                                                                                                                                                                                                                                            | 9 |
| enforce !walltime                                                                                                                                                                                                                                                                           | 9 |
| Enforce or don't enforce job "walltime" request. This is enforced by default.                                                                                                                                                                                                               | 9 |
| enforce pcpur                                                                                                                                                                                                                                                                               | 9 |
| enforce !pcpur                                                                                                                                                                                                                                                                              | 9 |
| Enforce or don't enforce job "pcpur" request. This is enforced by default.                                                                                                                                                                                                                  | 9 |
| enforce cput                                                                                                                                                                                                                                                                                | 9 |
| enforce !cput                                                                                                                                                                                                                                                                               | 9 |
| Enforce or don't enforce job "cput" request. This is enforced by default.                                                                                                                                                                                                                   | 9 |
| enforce cpupct                                                                                                                                                                                                                                                                              | 9 |
| enforce !cpupct                                                                                                                                                                                                                                                                             | 9 |
| Enforce or don't enforce job "cpupercent" request. This is not enforced by default.                                                                                                                                                                                                         | 9 |
| enforce file                                                                                                                                                                                                                                                                                | 9 |
| enforce !file                                                                                                                                                                                                                                                                               | 9 |
| Enforce or don't enforce job "file" request. This is enforced by default.                                                                                                                                                                                                                   | 9 |
| enforce hammer                                                                                                                                                                                                                                                                              | 9 |
| enforce !hammer                                                                                                                                                                                                                                                                             | 9 |
| Enforce or don't enforce the killing of processes of unauthorized users. This is not enforced by default.                                                                                                                                                                                   | 9 |
| enforce nokill                                                                                                                                                                                                                                                                              | 9 |
| enforce !nokill                                                                                                                                                                                                                                                                             | 9 |

|                                                                                                                                                                                                                  |   |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| Don't kill or kill the non-PBS processes if hammer code is enabled. This is don't kill by default.                                                                                                               | 9 |
| <code>enforce cpusets</code>                                                                                                                                                                                     | 9 |
| <code>enforce !cpusets</code>                                                                                                                                                                                    | 9 |
| Enforce or don't enforce cpusets. This is enforced by default.                                                                                                                                                   | 9 |
| <b>minnodemem</b>                                                                                                                                                                                                | 9 |
| The amount (in mb) of memory that must be present on a node before it can be considered for scheduling. Default is the smallest memory found on any node.                                                        | 9 |
| For example,                                                                                                                                                                                                     | 9 |
| <code>minnodemem 512mb</code>                                                                                                                                                                                    | 9 |
| <b>minnodecpus</b>                                                                                                                                                                                               | 9 |
| The number of working cpus that must be present on a node before it can be considered for scheduling. Default is the smallest number of CPUs found on any node.                                                  | 9 |
| For example,                                                                                                                                                                                                     | 9 |
| <code>minnodecpus 2</code>                                                                                                                                                                                       | 9 |
| <b>memreserved</b>                                                                                                                                                                                               | 9 |
| The amount of memory per node reserved for "system" overhead. The mom assumes that a job can depend upon starting up with (minnodemem - memreserved) mb of RAM available for user code per node. Default is 0MB. | 9 |
| For example,                                                                                                                                                                                                     | 9 |
| <code>memreserved 16mb</code>                                                                                                                                                                                    | 9 |
| <b>schd_quantum</b>                                                                                                                                                                                              | 9 |
| The minimum # of nodeboards (or ssinodes) to be assigned to a job. System defaults to 1.                                                                                                                         | 9 |
| For example,                                                                                                                                                                                                     | 9 |
| <code>schd_quantum 2</code>                                                                                                                                                                                      | 9 |
| The configuration file must be "secure". It must be owned by a user id and group id less than 10 and not be world writable.                                                                                      | 9 |
| <b>FILES</b>                                                                                                                                                                                                     | 9 |
| <code>\$PBS_SERVER_HOME/mom_priv</code>                                                                                                                                                                          | 9 |
| the default directory for configuration files, typical (/usr/spool/pbs)/mom_priv.                                                                                                                                | 9 |
| <code>\$PBS_SERVER_HOME/mom_logs</code>                                                                                                                                                                          | 9 |
| directory for log files recorded by the server.                                                                                                                                                                  | 9 |
| <code>\$PBS_SERVER_HOME/mom_priv/prologue</code>                                                                                                                                                                 | 9 |
| the administrative script to be run before job execution.                                                                                                                                                        | 9 |
| <code>\$PBS_SERVER_HOME/mom_priv/eiplogue</code>                                                                                                                                                                 | 9 |
| the administrative script to be run after job execution.                                                                                                                                                         | 9 |
| <b>Signal Handling</b>                                                                                                                                                                                           | 9 |
| Pbs_mom handles the following signals:                                                                                                                                                                           | 9 |
| <b>SIGHUP</b> causes pbs_mom to re-read its configuration file, close and reopen the log file, and reinitialize resource structures.                                                                             | 9 |
| <b>SIGALRM</b> results in a log file entry. The signal is used to limit the time taken by certain children processes, such as the prologue and epilogue.                                                         | 9 |
| <b>SIGINT and SIGTERM</b>                                                                                                                                                                                        | 9 |
| Result in pbs_mom terminating all running children and exiting. This is the ac-                                                                                                                                  | 9 |

|                                                                                                       |   |
|-------------------------------------------------------------------------------------------------------|---|
| tion for the following signals as well: SIGXCPU, SIGXFSZ, SIGCPULIM, and SIGSHUTDN.                   | 9 |
| SIGPIPE, SIGUSR1, SIGUSR2, SIGINFO                                                                    | 9 |
| are ignored.                                                                                          | 9 |
| All other signals have their default behavior installed.                                              | 9 |
| <b>EXIT STATUS</b>                                                                                    | 9 |
| If the mini-server command fails to begin operation, the server exits with a value greater than zero. | 9 |
| <b>SEE ALSO</b>                                                                                       | 9 |
| pbs_server(8B), pbs_scheduler_basl(8B), pbs_scheduler_tcl(8B)                                         | 9 |

**6.3. Initiate Execution Mini-server with Globus Support Operation** 9

**NAME** 9

`pbs_mom_globus` – start a pbs batch execution mini-server supporting Globus 9

**SYNOPSIS** 9

`pbs_mom_globus` [-c *config*] [-d *directory*] [-L *logfile*] [-M *MOM\_Globus\_port*] [-R *RPP\_Globus\_port*] [-r] [-x] 9

**DESCRIPTION** 9

The **pbs\_mom\_globus** command starts the operation of a batch **Machine Oriented Mini-server**, MOM supporting Globus, on the local host. Typically, this command will be in a local boot file such as */etc/rc.local*. To insure that the `pbs_mom_globus` command is not runnable by the general user community, the server will only execute if its real and effective uid is zero. 9

When `pbs_mom_globus` picks up a job for execution, the globus resource string *-/site=globus:<gatekeeper>* of the job is consulted and used to open up a connection to Globus. Pbs\_mom forks a process for the job, starts up a *globus-gass-server* on a 1 server per unique username scheme, transforms #PBS directive lines in the user's submission script into an RSL string and submits the job to Globus, and exiting out of the forked process. All Globus job state changes are communicated back to `pbs_mom_globus` through periodic polling. 9

When job fails to submit due to globus job initialization failures, or any non GRAM authentication failures, then error message gets dumped into stderr and user is sent email. 9

When job fails due to no user password, proxy credential from certificate, or credential has expired, or some sort of "handshaking" error, then user is sent email of the error, and job is placed on hold. 9

`Pbs_mom_globus` will record a diagnostic message in a log file for any error occurrence. The log files are maintained in the *mom\_globus\_logs* directory below the home directory of the server. If the log file cannot be opened, the diagnostic message is written to the system console. 9

**OPTIONS** 9

**-c config** Specify a alternative configuration file, see description below. If this is a relative file name it will be relative to `PBS_HOME/mom_globus_priv`, see the **-d** option. If the specified file cannot be opened, `pbs_mom_globus` will abort. If the **-c** option is not supplied, `pbs_mom_globus` will attempt to open the default 9

configuration file "config" in `PBS_HOME/mom_globus_priv`. If this file is not present, `pbs_mom_globus` will log the fact and continue. 9

**-d directory** Specifies the path of the directory which is the home of the servers working files, `PBS_HOME`. This option is typically used along with **-M** when debugging MOM Globus. The default directory is given by `$PBS_SERVER_HOME` which is typically `/usr/spool/PBS`. 9

**-L logfile** Specify an absolute path name for use as the log file. If not specified, MOM Globus will open a file named for the current date in the `PBS_HOME/mom_globus_logs` directory, see the **-d** option. 9

**-M port** Specifies the port number on which the mini-server with Globus will listen for batch requests. 9

**-R port** Specifies the port number on which the mini-server with Globus will listen for resource monitor requests. Both a UDP and a TCP port of this number will be used. 9

- r** Specifies the impact on jobs which were in execution when the mini-server shut down. With the **-r** option, MOM Globus will cancel submitted Globus jobs, mark the jobs as terminated, and notify the batch server which owns the job. 9  
 Normally the mini-server is started from the system boot file without the **-r** option. The mini-server will make no attempt to signal the former session of any job which may have been running when the mini-server terminated. It is assumed that on reboot, all processes have been killed. It will however attempt to cancel the Globus job. 9  
 If the **-r** option is used following a reboot, process ids (pids) may be reused and MOM may kill a process that is not a batch session. 9
- a alarm** Used to specify the alarm timeout in seconds for computing a resource. Every time a resource request is processed, an alarm is set for the given amount of time. If the request has not completed before the given time, an alarm signal is generated. The default is 5 seconds. 9
- x** Disables the check for privileged port resource monitor connections. This is used mainly for testing since the privileged port is the only mechanism used to prevent any ordinary user from connecting. 9

## CONFIGURATION FILE 9

The configuration file may be specified on the command line at program start with the **-c** flag. The use of this file is to provide several types of run time information to `pbs_mom_globus`: static resource names and values, external resources provided by a program to be run on request via a shell escape, and values to pass to internal set up functions at initialization (and re-initialization). 9

Each item type is on a single line with the component parts separated by white space. If the line starts with a hash mark (pound sign, #), the line is considered to be a comment and is skipped. 9

### Static Resources 9

For static resource names and values, the configuration file contains a list of resource names/values pairs, one pair per line and separated by white space. An Example of static resource names and values could be the number of tape drives of different types and could be specified by 9

```
tape3480      4 9
tape3420      2 9
tapedat       1 9
tape8mm       1 9
```

### Shell Commands 9

If the first character of the value is an exclamation mark (!), the entire rest of the line is saved to be executed through the services of the **system(3)** standard library routine. 9

The shell escape provides a means for the resource monitor to yield arbitrary information to the scheduler. Parameter substitution is done such that the value of any qualifier sent with the query, as explained below, replaces a token with a percent sign (%) followed by the name of the qualifier. For example, here is a configuration file line which gives a resource name of "escape": 9

```
escape      !echo %xxx %yyy 9
```

If a query for "escape" is sent with no qualifiers, the command executed would be "echo %xxx %yyy". If one qualifier is sent, "escape[xxx=hi there]", the command executed would be "echo hi there %yyy". If two qualifiers are sent, "escape[xxx=hi][yyy=there]", the command executed would be "echo hi there". If a qualifier is sent with no matching token in the command line, "escape[zzz=snafu]", an error is reported. 9

|                                                                                                                                                                                                                                                                                                              |   |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| Initialization Value                                                                                                                                                                                                                                                                                         | 9 |
| An initialization value directive has a name which starts with a dollar sign (\$) and must be known to MOM via an internal table. The entries in this table now are:                                                                                                                                         | 9 |
| clienthost                                                                                                                                                                                                                                                                                                   | 9 |
| which causes a host name to be added to the list of hosts which will be allowed to connect to MOM as long as they are using a privileged port. For example, here are two configuration file lines which will allow the hosts "fred" and "wilma" to connect:                                                  | 9 |
| \$clienthost        fred                                                                                                                                                                                                                                                                                     | 9 |
| \$clienthost        wilma                                                                                                                                                                                                                                                                                    | 9 |
| Two host name are always allowed to connection to pbs_mom_globus, "localhost" and the name returned to pbs_mom_globus by the system call gethostname(). These names need not be specified in the configuration file.                                                                                         | 9 |
| restricted                                                                                                                                                                                                                                                                                                   | 9 |
| which causes a host name to be added to the list of hosts which will be allowed to connect to MOM Globus without needing to use a privileged port. These names allow for wildcard matching. For example, here is a configuration file line which will allow queries from any host from the domain "ibm.com". | 9 |
| \$restricted        *.ibm.com                                                                                                                                                                                                                                                                                | 9 |
| The restriction which applies to these connections is that only internal queries may be made. No resources from a config file will be found. This is to prevent any shell commands from being run by a non-root process.                                                                                     | 9 |
| logevent                                                                                                                                                                                                                                                                                                     | 9 |
| which sets the mask that determines which event types are logged by pbs_mom_globus. For example:                                                                                                                                                                                                             | 9 |
| \$logevent 0x1fff                                                                                                                                                                                                                                                                                            | 9 |
| \$logevent 255                                                                                                                                                                                                                                                                                               | 9 |
| The first example would set the log event mask to 0x1ff (511) which enables logging of all events including debug events. The second example would set the mask to 0x0ff (255) which enables all events except debug events.                                                                                 | 9 |
| The configuration file must be "secure". It must be owned by a user id and group id less than 10 and not be world writable.                                                                                                                                                                                  | 9 |
| <b>FILES</b>                                                                                                                                                                                                                                                                                                 | 9 |
| \$PBS_SERVER_HOME/mom_globus_priv                                                                                                                                                                                                                                                                            | 9 |
| the default directory for configuration files, typical                                                                                                                                                                                                                                                       | 9 |
| (/usr/spool/pbs)/mom_globus_priv.                                                                                                                                                                                                                                                                            | 9 |
| \$PBS_SERVER_HOME/mom_globus_logs                                                                                                                                                                                                                                                                            | 9 |
| directory for log files recorded by the server.                                                                                                                                                                                                                                                              | 9 |
| \$PBS_SERVER_HOME/mom_globus_priv/prologue                                                                                                                                                                                                                                                                   | 9 |
| the administrative script to be run before job execution.                                                                                                                                                                                                                                                    | 9 |
| \$PBS_SERVER_HOME/mom_globus_priv/eiplogue                                                                                                                                                                                                                                                                   | 9 |
| the administrative script to be run after job execution.                                                                                                                                                                                                                                                     | 9 |
| <b>Signal Handling</b>                                                                                                                                                                                                                                                                                       | 9 |
| Pbs_mom_globus handles the following signals:                                                                                                                                                                                                                                                                | 9 |
| SIGHUP causes pbs_mom_globus to re-read its configuration file, close and reopen the log file, and reinitialize resource structures.                                                                                                                                                                         | 9 |



|                                                                                                                                                                          |   |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| SIGALRM results in a log file entry. The signal is used to limit the time taken by certain children processes, such as the prologue and epilogue.                        | 9 |
| SIGINT and SIGTERM                                                                                                                                                       | 9 |
| Result in pbs_mom_globus terminating all running children and exiting. This is the action for the following signals as well: SIGXCPU, SIGXFSZ, SIGCPULIM, and SIGSHUTDN. | 9 |
| SIGPIPE, SIGUSR1, SIGUSR2, SIGINFO                                                                                                                                       | 9 |
| are ignored.                                                                                                                                                             | 9 |
| All other signals have their default behavior installed.                                                                                                                 | 9 |
| <b>EXIT STATUS</b>                                                                                                                                                       | 9 |
| If the mini-server command fails to begin operation, the server exits with a value greater than zero.                                                                    | 9 |
| <b>SEE ALSO</b>                                                                                                                                                          | 9 |
| pbs_server(8B), pbs_scheduler_basl(8B), pbs_scheduler_tcl(8B)                                                                                                            | 9 |

|                                                                                                                                                                                                                                                                                                                          |   |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| <b>6.4. Disable Queue</b>                                                                                                                                                                                                                                                                                                | 9 |
| <b>NAME</b>                                                                                                                                                                                                                                                                                                              | 9 |
| qdisable – disable input to a pbs destination                                                                                                                                                                                                                                                                            | 9 |
| <b>SYNOPSIS</b>                                                                                                                                                                                                                                                                                                          | 9 |
| qdisable destination ...                                                                                                                                                                                                                                                                                                 | 9 |
| <b>DESCRIPTION</b>                                                                                                                                                                                                                                                                                                       | 9 |
| The <b>qdisable</b> command directs that a destination should no longer accept batch jobs.                                                                                                                                                                                                                               | 9 |
| The qdisable command sends a <i>Manage</i> request to the batch server specified by <i>destination</i> .                                                                                                                                                                                                                 | 9 |
| The enabled queue attribute is set to FALSE. If the batch request is accepted, the server will no longer accept <i>Queue Job</i> requests which specified the disabled queue. Jobs which already reside in the queue will continue to be processed. This allows a queue to be "drained."                                 | 9 |
| In order to execute qdisable, the user must have PBS Operation or Manager privilege.                                                                                                                                                                                                                                     | 9 |
| <b>OPERANDS</b>                                                                                                                                                                                                                                                                                                          | 9 |
| The qdisable command accepts one or more <i>destination</i> operands. The operands are one of three forms:                                                                                                                                                                                                               | 9 |
| queue                                                                                                                                                                                                                                                                                                                    | 9 |
| @server                                                                                                                                                                                                                                                                                                                  | 9 |
| queue@server                                                                                                                                                                                                                                                                                                             | 9 |
| If <i>queue</i> is specified, the request is to disable that queue at the default server. If the <i>@server</i> form is given, the request is to disable all the queues at that server. If a full destination identifier, <i>queue@server</i> , is given, the request is to disable the named queue at the named server. | 9 |
| <b>STANDARD ERROR</b>                                                                                                                                                                                                                                                                                                    | 9 |
| The qdisable command will write a diagnostic message to standard error for each error occurrence.                                                                                                                                                                                                                        | 9 |
| <b>EXIT STATUS</b>                                                                                                                                                                                                                                                                                                       | 9 |
| Upon successful processing of all the operands presented to the qdisable command, the exit status will be a value of zero.                                                                                                                                                                                               | 9 |
| If the qdisable command fails to process any operand, the command exits with a value greater than zero.                                                                                                                                                                                                                  | 9 |
| <b>SEE ALSO</b>                                                                                                                                                                                                                                                                                                          | 9 |
| pbs_server(8B), qmgr(8B), and qenable(8B)                                                                                                                                                                                                                                                                                | 9 |

|                                                                                                                                                                                                                                                                                                                       |   |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| <b>6.5. Enable Queue</b>                                                                                                                                                                                                                                                                                              | 9 |
| <b>NAME</b>                                                                                                                                                                                                                                                                                                           | 9 |
| qenable – enable input to a pbs destination                                                                                                                                                                                                                                                                           | 9 |
| <b>SYNOPSIS</b>                                                                                                                                                                                                                                                                                                       | 9 |
| qenable destination ...                                                                                                                                                                                                                                                                                               | 9 |
| <b>DESCRIPTION</b>                                                                                                                                                                                                                                                                                                    | 9 |
| The <b>qenable</b> command directs that a destination should accept batch jobs.                                                                                                                                                                                                                                       | 9 |
| The qenable command sends a <i>Manage</i> request to the batch server specified by <i>destination</i> .                                                                                                                                                                                                               | 9 |
| The enabled queue attribute is set to TRUE. If the request is accepted, that server will accept <i>Queue Job</i> requests which specified the queue.                                                                                                                                                                  | 9 |
| In order to execute qenable, the user must have PBS Operation or Manager privilege.                                                                                                                                                                                                                                   | 9 |
| <b>OPERANDS</b>                                                                                                                                                                                                                                                                                                       | 9 |
| The qenable command accepts one or more <i>destination</i> operands. The operands are one of three forms:                                                                                                                                                                                                             | 9 |
| queue                                                                                                                                                                                                                                                                                                                 | 9 |
| @server                                                                                                                                                                                                                                                                                                               | 9 |
| queue@server                                                                                                                                                                                                                                                                                                          | 9 |
| If <i>queue</i> is specified, the request is to enable that queue at the default server. If the <i>@server</i> form is given, the request is to enable all the queues at that server. If a full destination identifier, <i>queue@server</i> , is given, the request is to enable the named queue at the named server. | 9 |
| <b>STANDARD ERROR</b>                                                                                                                                                                                                                                                                                                 | 9 |
| The qenable command will write a diagnostic message to standard error for each error occurrence.                                                                                                                                                                                                                      | 9 |
| <b>EXIT STATUS</b>                                                                                                                                                                                                                                                                                                    | 9 |
| Upon successful processing of all the operands presented to the qenable command, the exit status will be a value of zero.                                                                                                                                                                                             | 9 |
| If the qenable command fails to process any operand, the command exits with a value greater than zero.                                                                                                                                                                                                                | 9 |
| <b>SEE ALSO</b>                                                                                                                                                                                                                                                                                                       | 9 |
| pbs_server(8B), qdisable(8B), and qmgr(8B)                                                                                                                                                                                                                                                                            | 9 |

|                                                                                                                                                                                      |                                                                                                                                                                                                                                                                                                       |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>6.6. Run Job</b>                                                                                                                                                                  | 9                                                                                                                                                                                                                                                                                                     |
| <b>NAME</b>                                                                                                                                                                          | 9                                                                                                                                                                                                                                                                                                     |
| qrun – run a pbs batch job                                                                                                                                                           | 9                                                                                                                                                                                                                                                                                                     |
| <b>SYNOPSIS</b>                                                                                                                                                                      | 9                                                                                                                                                                                                                                                                                                     |
| qrun [-H host][-a] job_identifier ...                                                                                                                                                | 9                                                                                                                                                                                                                                                                                                     |
| <b>DESCRIPTION</b>                                                                                                                                                                   | 9                                                                                                                                                                                                                                                                                                     |
| The <b>qrun</b> command is used to force a batch server to initiate the execution of a batch job. The job is run regardless of scheduling position, resource requirements, or state. | 9                                                                                                                                                                                                                                                                                                     |
| In order to execute qrun, the user must have PBS Operation or Manager privilege.                                                                                                     | 9                                                                                                                                                                                                                                                                                                     |
| <b>OPTIONS</b>                                                                                                                                                                       | 9                                                                                                                                                                                                                                                                                                     |
| -H host                                                                                                                                                                              | Specifies the host within the cluster on which the job(s) are to be run. The <i>host</i> argument is the name of a host that is a member of the cluster of hosts managed by the server. If the option is not specified, the server will select the "worst possible" host on which to execute the job. |
| -a                                                                                                                                                                                   | Specifies that the job is to be run "asynchronously".                                                                                                                                                                                                                                                 |
| <b>OPERANDS</b>                                                                                                                                                                      | 9                                                                                                                                                                                                                                                                                                     |
| The qrun command accepts one or more <i>job_identifier</i> operands of the form:                                                                                                     | 9                                                                                                                                                                                                                                                                                                     |
| sequence_number[.server_name][@server]                                                                                                                                               | 9                                                                                                                                                                                                                                                                                                     |
| See the description under "Job Identifier" in section 2.7.6 in the ERS.                                                                                                              | 9                                                                                                                                                                                                                                                                                                     |
| <b>STANDARD ERROR</b>                                                                                                                                                                | 9                                                                                                                                                                                                                                                                                                     |
| The qrun command will write a diagnostic message to standard error for each error occurrence.                                                                                        | 9                                                                                                                                                                                                                                                                                                     |
| <b>EXIT STATUS</b>                                                                                                                                                                   | 9                                                                                                                                                                                                                                                                                                     |
| Upon successful processing of all the operands presented to the qrun command, the exit status will be a value of zero.                                                               | 9                                                                                                                                                                                                                                                                                                     |
| If the qrun command fails to process any operand, the command exits with a value greater than zero.                                                                                  | 9                                                                                                                                                                                                                                                                                                     |
| <b>SEE ALSO</b>                                                                                                                                                                      | 9                                                                                                                                                                                                                                                                                                     |
| pbs_server(8B), pbs_sched(8B), qmgr(8B), pbs_runjob(3B)                                                                                                                              | 9                                                                                                                                                                                                                                                                                                     |

|                                                                                                                                                                                                                                                                                                         |   |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| <b>6.7. Start Queue</b>                                                                                                                                                                                                                                                                                 | 9 |
| <b>NAME</b>                                                                                                                                                                                                                                                                                             | 9 |
| qstart – start pbs batch job processing at a destination                                                                                                                                                                                                                                                | 9 |
| <b>SYNOPSIS</b>                                                                                                                                                                                                                                                                                         | 9 |
| qstart destination ...                                                                                                                                                                                                                                                                                  | 9 |
| <b>DESCRIPTION</b>                                                                                                                                                                                                                                                                                      | 9 |
| The <b>qstart</b> command directs that a destination should process batch jobs. If the destination is an execution queue, the server will begin to schedule jobs that reside in the queue for execution. If the destination is a routing queue, the server will begin to route jobs from that queue.    | 9 |
| The qstart command sends a <i>Manage</i> request to the batch server specified by <i>destination</i> . The started queue attribute is set to TRUE.                                                                                                                                                      | 9 |
| In order to execute qstart, the user must have PBS Operation or Manager privilege.                                                                                                                                                                                                                      | 9 |
| <b>OPERANDS</b>                                                                                                                                                                                                                                                                                         | 9 |
| The qstart command accepts one or more <i>destination</i> operands. The operands are one of three forms:                                                                                                                                                                                                | 9 |
| queue                                                                                                                                                                                                                                                                                                   | 9 |
| @server                                                                                                                                                                                                                                                                                                 | 9 |
| queue@server                                                                                                                                                                                                                                                                                            | 9 |
| If <i>queue</i> is specified, the request is to start that queue at the default server. If the @server form is given, the request is to start all queues at that server. If a full destination identifier, <i>queue@server</i> , is given, the request is to start the named queue at the named server. | 9 |
| <b>STANDARD ERROR</b>                                                                                                                                                                                                                                                                                   | 9 |
| The qstart command will write a diagnostic message to standard error for each error occurrence.                                                                                                                                                                                                         | 9 |
| <b>EXIT STATUS</b>                                                                                                                                                                                                                                                                                      | 9 |
| Upon successful processing of all the operands presented to the qstart command, the exit status will be a value of zero.                                                                                                                                                                                | 9 |
| If the qstart command fails to process any operand, the command exits with a value greater than zero.                                                                                                                                                                                                   | 9 |
| <b>SEE ALSO</b>                                                                                                                                                                                                                                                                                         | 9 |
| pbs_server(8B), qstop(8B), and qmgr(8B)                                                                                                                                                                                                                                                                 | 9 |

|                                                                                                                                                                                                                                                                                                          |   |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| <b>6.8. Stop Queue</b>                                                                                                                                                                                                                                                                                   | 9 |
| <b>NAME</b>                                                                                                                                                                                                                                                                                              | 9 |
| qstop – stop pbs batch job processing at a destination                                                                                                                                                                                                                                                   | 9 |
| <b>SYNOPSIS</b>                                                                                                                                                                                                                                                                                          | 9 |
| qstop destination ...                                                                                                                                                                                                                                                                                    | 9 |
| <b>DESCRIPTION</b>                                                                                                                                                                                                                                                                                       | 9 |
| The <b>qstop</b> command directs that a destination should stop processing batch jobs. If the destination is a execution queue, the server will cease scheduling jobs that reside in the queue for execution. If the destination is a routing queue, the server will cease routing jobs from that queue. | 9 |
| The qstop command sends a <i>Manage</i> request to the batch server specified by <i>destination</i> . The started queue attribute is set to FALSE.                                                                                                                                                       | 9 |
| In order to execute qstop, the user must have PBS Operation or Manager privilege.                                                                                                                                                                                                                        | 9 |
| <b>OPERANDS</b>                                                                                                                                                                                                                                                                                          | 9 |
| The qstop command accepts one or more <i>destination</i> operands. The operands are one of three forms:                                                                                                                                                                                                  | 9 |
| queue                                                                                                                                                                                                                                                                                                    | 9 |
| @server                                                                                                                                                                                                                                                                                                  | 9 |
| queue@server                                                                                                                                                                                                                                                                                             | 9 |
| If <i>queue</i> is specified, the request is to stop that queue at the default server. If the @server form is given, the request is to stop all the queues at that server. If a full destination identifier, <i>queue@server</i> , is given, the request is to stop the named queue at the named server. | 9 |
| <b>STANDARD ERROR</b>                                                                                                                                                                                                                                                                                    | 9 |
| The qstop command will write a diagnostic message to standard error for each error occurrence.                                                                                                                                                                                                           | 9 |
| <b>EXIT STATUS</b>                                                                                                                                                                                                                                                                                       | 9 |
| Upon successful processing of all the operands presented to the qstop command, the exit status will be a value of zero.                                                                                                                                                                                  | 9 |
| If the qstop command fails to process any operand, the command exits with a value greater than zero.                                                                                                                                                                                                     | 9 |
| <b>SEE ALSO</b>                                                                                                                                                                                                                                                                                          | 9 |
| pbs_server(8B), qstart(8B), and qmgr(8B)                                                                                                                                                                                                                                                                 | 9 |

**6.9. Terminate Server Operation** 9**NAME** 9

qterm – terminate processing by a pbs batch server 9

**SYNOPSIS** 9

qterm [-t type] [-m] [-s] [server...] 9

**DESCRIPTION** 9

The **qterm** command terminates a batch server. The command sends a *Server Shutdown* batch request to the server. When a server receives a Shutdown request, The server will go into *Terminating* state. No new jobs will be allowed to be started into execution nor enqueued into the server. The impact on jobs currently being run by the server depends on the type of shut down requested as described below. The qterm command will not exit until the server has completed its shutdown procedure.

In order to execute qterm, the user must have PBS Operation or Manager privilege. 9

**OPTIONS** 9

-m Specifies that all known pbs\_mom daemons should also be told to shut down. This request is relayed by the Server to each Mom. 9

-s Specifies that the Job Scheduler, pbs\_sched, should also be terminated. 9

-t type Specifies the type of shut down. The types are: 9

immediate 9

All running jobs are to immediately stop execution. The value passed in the *Shutdown* request is {SHUT\_IMMEDIATE}. 9

If checkpoint is supported, running jobs that can be checkpointed are checkpointed, terminated, and requeued. If checkpoint is not supported or the job cannot be checkpointed, running jobs are requeued if the rerunable attribute is true. Otherwise, jobs are killed. 9

Normally the server will not shutdown until there are no jobs in the running state. If the server is unable to contact the MOM of running job, the job is still listed as running. The server may be forced down by a second “qterm -t immediate” command. 9

delay If checkpoint is supported, running jobs that can be checkpointed are checkpointed, terminated, and requeued. If a job cannot be checkpointed, but can be rerun, the job is terminated and requeued. Otherwise, running jobs are allowed to continue to run. Note, the operator or administrator may use the qrerun and qdel commands to remove running jobs. The value passed in the *Shutdown* request is {SHUT\_DELAY}. 9

quick This is the default action if the -t option is not specified. This option is used when you wish that running jobs be left running when the server shuts down. The server will cleanly shutdown and can be restarted when desired. Upon restart of the server, jobs that continue to run are shown as running; jobs that terminated during the server’s absence will be placed into the exiting state. The value passed in the *Shutdown* request is 9

**OPERANDS** 9

The *server* operand specifies which servers are to shutdown. If no servers are given, then the default server will be terminated. 9

**STANDARD ERROR** 9

The qterm command will write a diagnostic message to standard error for each error occurrence. 9  
9

**EXIT STATUS** 9

Upon successful processing of all the operands presented to the qterm command, the exit status will be a value of zero. 9  
9

If the qterm command fails to process any operand, the command exits with a value greater than zero. 9  
9

**SEE ALSO** 9

pbs\_server(8B), qmgr(8B), pbs\_resources\_aix4(7B), pbs\_resources\_iris5(7B), pbs\_re- 9  
sources\_sp2(7B), pbs\_resources\_sunos4(7B), and pbs\_resources\_unicos8(7B) 9





## 7. Administrator Commands 9

A batch administrator is a user of the system who is granted the highest level of privilege in the batch system. The batch administrator is able to perform all operator functions and additionally modify queue and server configurations. 9

In addition to the general user commands available to an operator with increased privilege and the special operator commands, the administrator has access to the **qmgr** command. Information on the Resource Monitor and Schedulers are also provided for the administrator. 9

There are three different schedulers provided with PBS. The yacc/lex based scheduler is driven by a procedure written in a "C - like language." This BATCH Scheduler Language (BASL) scheduler is sometimes referred to (incorrectly) as the "rules" based scheduler. 9

Another scheduler based on Tcl<sup>1</sup> another is a C language framework. is provided for sites which wish a different approach. The third scheduler is a framework (main program) written in the C language. The site must supply the primary function (in C) which implements the scheduling policy. 9

---

<sup>1</sup> Tool Control Langague - see "Tcl and the Tk Toolkit" by John K. Ousterhout.

## 7.1. PBS Job Schedulers 9

### 7.1.1. C Based Job Scheduler 9

#### NAME 9

`pbs_sched_cc` – pbs C scheduler 9

#### SYNOPSIS 9

`pbs_sched` [-a alarm] [-d home] [-L logfile] [-p file] [-S port] [-R port] [-c file] 9

#### DESCRIPTION 9

The **pbs\_sched** program runs in conjunction with the PBS server. It queries the server about the state of PBS and communicates with **pbs\_resmon** to get information about the status of running jobs, memory available etc. It then makes decisions as to what jobs to run. 9

`pbs_sched` must be executed with root permission. 9

#### OPTIONS 9

**-a alarm** This specifies the time in seconds to wait for a schedule run to finish. If a script takes too long to finish, an alarm signal is sent, and the scheduler is restarted. If a core file does not exist in the current directory, **abort()** is called and a core file is generated. The default for *alarm* is 180 seconds. 9

**-d home** This specifies the PBS home directory, `PBS_HOME`. The current working directory of the scheduler is `PBS_HOME/sched_priv`. If this option is not given, `PBS_HOME` defaults to `$PBS_SERVER_HOME` as defined during the PBS build procedure. 9

**-L logfile** Specifies an absolute path name of the file to use as the log file. If not specified, the scheduler will open a file named for the current date in the `PBS_HOME/sched_logs` directory (see the `-d` option). 9

**-p file** This specifies the "print" file. Any output from the C code which is written to standard out or standard error will be written to this file. If this option is not given, the file used will be `PBS_HOME/sched_priv/sched_out`. See the `-d` option. 9

**-S port** This specifies the port to use. If this option is not given, the default port for the PBS scheduler is used. 9

**-R port** This specifies the resource monitor port to use. If this option is not given, the default port for the PBS mom is used. NOTE: this option only makes the mom port available to the scheduler writer. It doesn't force them to use it. 9

**-c file** Specify a configuration file, see description below. If this is a relative file name it will be relative to `PBS_HOME/sched_priv`, see the `-d` option. If the `-c` option is not supplied, `pbs_sched` will not attempt to open a configuration file. 9

The options that specify file names may be absolute or relative. If they are relative, their root directory will be `PBS_HOME/sched_priv`. 9

#### USAGE 9

This version of the scheduler requires knowledge of the C language and the PBS API. Source code is provided for a main program for the scheduler. The site must supply the heart of the program. When invoked, the main program performs general initialization and housekeeping chores. Then a locally supplied function, *schedinit()* is called to perform site specific initialization. 9

In the main loop, a locally supplied function, *schedule()* is called to make the scheduling decisions and perform any required actions. Information about jobs and queues is obtained from the Server through the standard PBS API as found in libifl.a. Information about the execution host(s) is obtained from the Resource Monitor. Routines to communicate with the Resource Monitor are found in libnet.a.

If the processing takes more than the allotted time, the scheduler will restart itself. The default amount of time is three minutes. This can be changed with the -a option.

On receipt of a SIGHUP signal, the scheduler will close and reopen its log file and reread its configuration file (if any).

**CONFIGURATION FILE**

A configuration file may be specified with the -c option. This file may be used to specify the hosts (servers) which are allowed to connect to pbs\_sched. The hosts are specified in the configuration file in a manor identical to that used in pbs\_mom. There is one line per host with the syntax:

\$clienthost hostname

where clienthost and hostname are separated by white space.

Two host names are always allowed to connection to pbs\_sched, "localhost" and the name returned to pbs\_sched by the system call gethostname(). These names need not be specified in the configuration file.

The configuration file must be "secure". It must be owned by a user id and group id less than 10 and not be world writable.

**FILES**

\$PBS\_SERVER\_HOME/sched\_priv  
the default directory for configuration files, typically (/usr/spool/pbs)/sched\_priv.

**Signal Handling**

A C based scheduler will handle the following signals:

**SIGHUP**

The server will close and reopen its log file and reread the config file if one exists.

**SIGALRM**

If the site supplied scheduling module exceeds the time limit, the Alarm will cause the scheduler to attempt to core dump and restart itself.

**SIGINT and SIGTERM**

Will result in an orderly shutdown of the scheduler.

All other signals have the default action installed.

**EXIT STATUS**

Upon normal termination, an exit status of zero is returned.

**SEE ALSO**

pbs\_sched\_rule(8B), pbs\_sched\_tcl(8B), pbs\_server(8B), pbs\_mom(8B)

## 7.1.2. Yacc/lex Based Job Scheduler 9

**NAME** 9

pbs\_sched\_basl – pbs BASL scheduler 9

**SYNOPSIS** 9

pbs\_sched [-d home] [-L logfile] [-p print\_file] [-a alarm] [-S port] [-c configfile] 9

**DESCRIPTION** 9

The **pbs\_sched** command starts the operation of a batch scheduler on the local host. It runs in conjunction with the PBS server. It queries the server about the state of PBS and communicates with **pbs\_mom** to get information about the status of running jobs, memory available etc. It then makes decisions as to what jobs to run. 9

Typically, this command will be in a local boot file such as */etc/rc.local*. 9

pbs\_sched must be executed with root permission. 9

**OPTIONS** 9

-d home 9

Specifies the name of the PBS home directory, PBS\_HOME. If not specified, the value of \$PBS\_SERVER\_HOME as defined at compile time is used. Also see the -L option. 9

-L logfile 9

Specifies an absolute path name of the file to use as the log file. If not specified, the scheduler will open a file named for the current date in the PBS\_HOME/sched\_logs directory. See the -d option. 9

-p print\_file 9

This specifies the "print" file. Any output from the scheduler code which is written to standard out or standard error will be written to this file. If this option is not given, the file used will be \$PBS\_HOME/sched\_priv/sched\_out. See the -d option. 9

-a alarm 9

This specifies the time in seconds to wait for a schedule run to finish. If a scheduling iteration takes too long to finish, an alarm signal is sent, and the scheduler is restarted. If a core file does not exist in the current directory, abort() is called and a core file is generated. The default for alarm is 180 seconds. 9

-S port 9

Specifies a port on which to talk to the server. This option is not required. It merely overrides the default PBS scheduler port. 9

-c configfile 9

Specify a configuration file, see description below. If this is a relative file name it will be relative to PBS\_HOME/sched\_priv, see the -d option. If the -c option is not supplied, pbs\_sched will not attempt to open a configuration file. In BASL, this config file is almost always needed because it is where the list of servers, nodes, and host resource queries are specified by the administrator. 9

**USAGE** 9

This version of the scheduler requires knowledge of the BASL language. The site must first write a function called *sched\_main()* (and all functions supporting it) using BASL constructs, and then translate the functions into C using the BASL compiler *basl2c*, which would also attach a main program to the resulting code. This main program performs general initialization and housekeeping chores such as setting up local socket to communicate with the server running on the same machine, cd-ing to the priv directory, opening log files, opening configuration file (if any), setting up locks, forking the child to become a daemon, initializing a 9

scheduling cycle (i.e. get node attributes that are static in nature), setting up the signal handlers, executing global initialization assignment statements specified by the scheduler writer, and finally sitting on a loop waiting for a scheduling command from the server. When the server sends the scheduler an appropriate scheduling command {SCH\_SCHEDULE\_NEW, SCH\_SCHEDULE\_TERM, SCH\_SCHEDULE\_TIME, SCH\_SCHEDULE\_RECYC, SCH\_SCHEDULE\_CMD, SCH\_SCHEDULE\_FIRST}, information about server(s), jobs, queues, and execution host(s) are obtained, and then *sched\_main()* is called.

**SCHEDULING LANGUAGE**

The BAtch Scheduling Language (BASL) is a C-like procedural language. It provides a number of constructs and predefined functions that facilitate dealing with scheduling issues. Information about a PBS server, the queues that it owns, jobs residing on each queue, and the computational nodes where jobs can be run, are accessed via the BASL data types Server, Que, Job, CNode, Set Server, Set Que, Set Job, and Set CNode.

The following simple *sched\_main()* will cause the server to run all queued jobs on the local server:

```

sched_main()
{
    Server  s;
    Que     q;
    Job     j;
    Set Que queues;
    Set Job jobs;

    s = AllServersLocalHostGet(); // get local server
    queues = ServerQueuesGet(s);

    foreach( q in queues ) {
        jobs = QueJobsGet(q);
        foreach( j in jobs ) {
            JobAction(j, SYNCRUN, NULLSTR);
        }
    }
}

```

For a more complete discussion of the Batch Scheduler Language, see **basl2c(1B)**.

**CONFIGURATION FILE**

A configuration file may be specified with the -c option. This file is used to specify the (1) hosts which are allowed to connect to pbs\_sched, (2) the list of server hosts for which the scheduler writer wishes the system to periodically check for status, queues, and jobs info, (3) list of execution hosts for which the scheduler writer wants the system to periodically check for information like state, property, and so on, and (4) various queries to send to each execution host.

(1) specifying client hosts:

The hosts allowed to connect to pbs\_sched are specified in the configuration file in a manner identical to that used in pbs\_mom. There is one line per host using the syntax:

```
$clienthost  hostname
```

where *clienthost* and *hostname* are separated by white space. Two host names are always allowed to connection to pbs\_sched: "localhost" and the name returned to

|                                                                                                                                                                                                                                                           |   |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| pbs_sched by the system call gethostname(). These names need not be specified in the configuration file.                                                                                                                                                  | 9 |
| (2) specifying list of servers:                                                                                                                                                                                                                           | 9 |
| The list of servers is specified in a one host per line manner, using the syntax:                                                                                                                                                                         | 9 |
| <pre>\$serverhost hostname port_number</pre>                                                                                                                                                                                                              | 9 |
| or where <code>\$server_host</code> , <code>hostname</code> , and <code>port_number</code> are separated by white space.                                                                                                                                  | 9 |
| If <code>port_number</code> is 0, then the default PBS server port will be used.                                                                                                                                                                          | 9 |
| Regardless of what has been specified in the file, the list of servers will always include the local server - one running on the same host where the scheduler is running.                                                                                | 9 |
| Within the BASL code, access to data of the list of servers is done by calling <code>AllServersGet()</code> , or <code>AllServersLocalHostGet()</code> which returns the local server on the list.                                                        | 9 |
| (3) specifying the list of execution hosts:                                                                                                                                                                                                               | 9 |
| The list of execution hosts (nodes), whose MOMs are to be queried from the scheduler, is specified in a one host per line manner, using the syntax:                                                                                                       | 9 |
| <pre>\$momhost hostname port_number</pre>                                                                                                                                                                                                                 | 9 |
| where <code>\$momhost</code> , <code>hostname</code> , and <code>port_number</code> are separated by white space.                                                                                                                                         | 9 |
| If <code>port_number</code> is 0, then the default PBS MOM port will be used.                                                                                                                                                                             | 9 |
| The BASL function <code>AllNodesGet()</code> , or <code>ServerNodesGet(AllServersLocalHostGet())</code> is available for getting the list of nodes known to the local system.                                                                             | 9 |
| (4) specifying the list of host resources:                                                                                                                                                                                                                | 9 |
| For specifying the list of host resource queries to send to each execution host's MOM, the following syntax is used:                                                                                                                                      | 9 |
| <pre>\$node node_name CNode..Get host_resource</pre>                                                                                                                                                                                                      | 9 |
| <code>node_name</code> should be the same hostname string that was specified in a <code>\$momhost</code> line. A <code>node_name</code> value of "*" (wildcard) means to match any node.                                                                  | 9 |
| Please consult section 9 of the PBS ERS (Resource Monitor/Resources) for a list of possible values to <code>host_resource</code> parameter.                                                                                                               | 9 |
| <code>CNode..Get</code> refers to the actual function name that is called from the scheduler code to obtain the return values to host resource queries. The list of <code>CNode..Get</code> function names that can appear in the configuration file are: | 9 |

|                                        |   |
|----------------------------------------|---|
| <b>STATIC:</b>                         | 9 |
| =====                                  | 9 |
| <b>CNodePropertiesGet</b>              | 9 |
| <b>CNodeVendorGet</b>                  | 9 |
| <b>CNodeNumCpusGet</b>                 | 9 |
| <b>CNodeOsGet</b>                      | 9 |
| <b>CNodeMemTotalGet[type]</b>          | 9 |
| <b>CNodeNetworkBwGet[type]</b>         | 9 |
| <b>CNodeSwapSpaceTotalGet[name]</b>    | 9 |
| <b>CNodeDiskSpaceTotalGet[name]</b>    | 9 |
| <b>CNodeDiskInBwGet[name]</b>          | 9 |
| <b>CNodeDiskOutBwGet[name]</b>         | 9 |
| <b>CNodeTapeSpaceTotalGet[name]</b>    | 9 |
| <b>CNodeTapeInBwGet[name]</b>          | 9 |
| <b>CNodeTapeOutBwGet[name]</b>         | 9 |
| <b>CNodeSrfsSpaceTotalGet[name]</b>    | 9 |
| <b>CNodeSrfsInBwGet[name]</b>          | 9 |
| <b>CNodeSrfsOutBwGet[name]</b>         | 9 |
| <br>                                   |   |
| <b>DYNAMIC:</b>                        | 9 |
| =====                                  | 9 |
| <b>CNodeIdleTimeGet</b>                | 9 |
| <b>CNodeLoadAveGet</b>                 | 9 |
| <b>CNodeMemAvailGet[type]</b>          | 9 |
| <b>CNodeSwapSpaceAvailGet[name]</b>    | 9 |
| <b>CNodeSwapInBwGet[name]</b>          | 9 |
| <b>CNodeSwapOutBwGet[name]</b>         | 9 |
| <b>CNodeDiskSpaceReservedGet[name]</b> | 9 |
| <b>CNodeDiskSpaceAvailGet[name]</b>    | 9 |
| <b>CNodeTapeSpaceAvailGet[name]</b>    | 9 |
| <b>CNodeSrfsSpaceReservedGet[name]</b> | 9 |
| <b>CNodeSrfsSpaceAvailGet[name]</b>    | 9 |
| <b>CNodeCpuPercentIdleGet</b>          | 9 |
| <b>CNodeCpuPercentSysGet</b>           | 9 |
| <b>CNodeCpuPercentUserGet</b>          | 9 |
| <b>CNodeCpuPercentGuestGet</b>         | 9 |

STATIC function names return values that are obtained only during the first scheduling cycle, or when the scheduler is instructed to reconfig; whereas, DYNAMIC function names return attribute values that are taken at every subsequent scheduling cycle. 9

**name** and **type** are arbitrarily defined. For example, you can choose to have **name** defined as "\$FASTDIR" for the CNodeSrfs\* calls, and a sample configuration file entry would look like: 9

```

$node unicos8 CNodeSrfsSpaceAvailGet[$FASTDIR]          9
                quota[type=ares_avail,dir=$FASTDIR]      9
    
```

So in a BASL code, if you call CNodeSrfsSpaceAvailGet(node, "\$FASTDIR"), then it will return the value to the query "quota[type=ares\_avail,dir=\$FASTDIR]" (3rd parameter) as sent to the node's MOM. 9

By default, the scheduler has already internally defined the following mappings, which can be overridden in the configuration file: 9



| keyword | node_name | CNode..Get       | host_resource | 9 |
|---------|-----------|------------------|---------------|---|
| =====   | =====     | =====            | =====         | 9 |
| \$node  | *         | CNodeOsGet       | arch          | 9 |
| \$node  | *         | CNodeLoadAveGet  | loadave       | 9 |
| \$node  | *         | CNodeIdleTimeGet | idletime      | 9 |

The above means that for all declared nodes (via \$momhost), the host queries arch, loadave, and idletime will be sent to each node's MOM. The value to arch is obtained internally by the system during the first scheduling cycle because it falls under STATIC category, while values to loadave and idletime are taken at every scheduling iteration because they fall under the DYNAMIC category. Access to the return values is done by calling *CNodeOsGet(node)*, *CNodeLoadAveGet(node)*, and *CNodeIdleTimeGet(node)*, respectively. The following are some sample \$node arguments that you may put in the configuration file.

|                     |                                   |                                              |             |
|---------------------|-----------------------------------|----------------------------------------------|-------------|
| node_name           | CNode..Get                        | host res                                     | 9           |
| =====               | =====                             | =====                                        | 9           |
| <sunos4_nodename>   | CNodeIdleTimeGet                  | idletime                                     | 9           |
| <sunos4_nodename>   | CNodeLoadAveGet                   | loadave                                      | 9           |
| <sunos4_nodename>   | CNodeMemTotalGet[real]            | physmem                                      | 9           |
| <sunos4_nodename>   | CNodeMemTotalGet[virtual]         | totmem                                       | 9           |
| <sunos4_nodename>   | CNodeMemAvailGet[virtual]         | availmem                                     | 9           |
| <irix5_nodename>    | CNodeNumCpusGet                   | ncpus                                        | 9           |
| <irix5_nodename>    | CNodeMemTotalGet[real]            | physmem                                      | 9           |
| <irix5_nodename>    | CNodeMemTotalGet[virtual]         | totmem                                       | 9           |
| <irix5_nodename>    | CNodeIdleTimeGet                  | idletime                                     | 9           |
| <irix5_nodename>    | CNodeLoadAveGet                   | loadave                                      | 9           |
| <irix5_nodename>    | CNodeMemAvailGet[virtual]         | availmem                                     | 9           |
| <linux_nodename>    | CNodeNumCpusGet                   | ncpus                                        | 9           |
| <linux_nodename>    | CNodeMemTotalGet[real]            | physmem                                      | 9           |
| <linux_nodename>    | CNodeMemTotalGet[virtual]         | totmem                                       | 9           |
| <linux_nodename>    | CNodeIdleTimeGet                  | idletime                                     | 9           |
| <linux_nodename>    | CNodeLoadAveGet                   | loadave                                      | 9           |
| <linux_nodename>    | CNodeMemAvailGet[virtual]         | availmem                                     | 9           |
| <solaris5_nodename> | CNodeIdleTimeGet                  | idletime                                     | 9           |
| <solaris5_nodename> | CNodeLoadAveGet                   | loadave                                      | 9           |
| <solaris5_nodename> | CNodeNumCpusGet                   | ncpus                                        | 9           |
| <solaris5_nodename> | CNodeMemTotalGet[real]            | physmem                                      | 9           |
| <aix4_nodename>     | CNodeIdleTimeGet                  | idletime                                     | 9           |
| <aix4_nodename>     | CNodeLoadAveGet                   | loadave                                      | 9           |
| <aix4_nodename>     | CNodeMemTotalGet[virtual]         | totmem                                       | 9           |
| <aix4_nodename>     | CNodeMemAvailGet[virtual]         | availmem                                     | 9           |
| <unicos8_nodename>  | CNodeIdleTimeGet                  | idletime                                     | 9           |
| <unicos8_nodename>  | CNodeLoadAveGet                   | loadave                                      | 9           |
| <unicos8_nodename>  | CNodeNumCpusGet                   | ncpus                                        | 9           |
| <unicos8_nodename>  | CNodeMemTotalGet[real]            | physme                                       | 9           |
| <unicos8_nodename>  | CNodeMemAvailGet[virtual]         | availmem                                     | 9           |
| <unicos8_nodename>  | CNodeSwapSpaceTotalGet[primary]   | swaptotal                                    | 9           |
| <unicos8_nodename>  | CNodeSwapSpaceAvailGet[primary]   | swapavail                                    | 9           |
| <unicos8_nodename>  | CNodeSwapInBwGet[primary]         | swpinrate                                    | 9           |
| <unicos8_nodename>  | CNodeSwapOutBwGet[primary]        | swpoutrate                                   | 9           |
| <unicos8_nodename>  | CNodePercentIdleGet               | cpuidle                                      | 9           |
| <unicos8_nodename>  | CNodePercentSysGet                | cpuunix                                      | 9           |
| <unicos8_nodename>  | CNodePercentGuestGet              | cpuguest                                     | 9           |
| <unicos8_nodename>  | CNodePercentUusrGet               | cpuuser                                      | 9           |
| <unicos8_nodename>  | CNodeSrfsSpaceAvailGet[\$FASTDIR] | quota[type<br>=ares_avail,<br>dir=\$FASTDIR] | 9<br>9<br>9 |
| <unicos8_nodename>  | CNodeSrfsSpaceAvailGet[\$BIGDIR]  | quota[type<br>=ares_avail,<br>dir=\$BIGDIR]  | 9<br>9<br>9 |
| <unicos8_nodename>  | CNodeSrfsSpaceAvailGet[\$WRKDIR]  | quota[type                                   | 9           |

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |                                                                                                                                               |               |   |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|---------------|---|
|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |                                                                                                                                               | =ares_avail,  | 9 |
|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |                                                                                                                                               | dir=\$WRKDIR] | 9 |
| <sp2_nodename>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | CNodeLoadAveGet                                                                                                                               | loadave       | 9 |
| <p>Suppose you have an execution host that is of irix5 os type, then the &lt;irix5_node_name&gt; entries will be consulted by the scheduler. The initial scheduling cycle would involve sending the STATIC queries ncpus, physmem, totmem to the execution host's MOM, and access to return values of the queries is done via CNodeNumCpusGet(node), CNodeMemTotalGet(node, "real"), CNodeMemTotalGet(node, "virtual") respectively, where node is the CNode representation of the execution host. The subsequent scheduling cycles will only send DYNAMIC queries idletime, loadave, and availmem, and access to the return values of the queries is done via CNodeIdleTimeGet(node), CNodeLoadAveGet(node), CNodeMemAvailGet(node, "virtual"). respectively.</p> |                                                                                                                                               |               |   |
| <p>"Later" entries in the config file take precedence.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |                                                                                                                                               |               |   |
| <p>The configuration file must be "secure". It must be owned by a user id and group id less than 10 and not be world writable.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |                                                                                                                                               |               |   |
| <p>On receipt of a SIGHUP signal, the scheduler will close and reopen its log file and reread its configuration file (if any).</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |                                                                                                                                               |               |   |
| <p><b>FILES</b></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |                                                                                                                                               |               |   |
| \$PBS_SERVER_HOME/sched_priv                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |                                                                                                                                               |               | 9 |
|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | the default directory for configuration files, typically (/usr/spool/pbs)/sched_priv.                                                         |               | 9 |
| <p><b>Signal Handling</b></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |                                                                                                                                               |               |   |
| <p>A C based scheduler will handle the following signals:</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |                                                                                                                                               |               |   |
| <p><b>SIGHUP</b></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |                                                                                                                                               |               |   |
|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | The server will close and reopen its log file and reread the config file if one exists.                                                       |               | 9 |
| <p><b>SIGALRM</b></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |                                                                                                                                               |               |   |
|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | If the site supplied scheduling module exceeds the time limit, the Alarm will cause the scheduler to attempt to core dump and restart itself. |               | 9 |
| <p><b>SIGINT and SIGTERM</b></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |                                                                                                                                               |               |   |
|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Will result in an orderly shutdown of the scheduler.                                                                                          |               | 9 |
| <p>All other signals have the default action installed.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |                                                                                                                                               |               |   |
| <p><b>EXIT STATUS</b></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |                                                                                                                                               |               |   |
|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Upon normal termination, an exit status of zero is returned.                                                                                  |               | 9 |
| <p><b>SEE ALSO</b></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |                                                                                                                                               |               |   |
|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | basl2c(1B), pbs_sched_tcl(8B), pbs_server(8B), pbs_mom(8B)                                                                                    |               | 9 |

|                                                                                                                                                                                                                                                                                    |                                                                                                                                                                                                                                                                                                                                     |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 7.1.3. Tcl Based Job Scheduler                                                                                                                                                                                                                                                     | 9                                                                                                                                                                                                                                                                                                                                   |
| <b>NAME</b>                                                                                                                                                                                                                                                                        | 9                                                                                                                                                                                                                                                                                                                                   |
| pbs_sched_tcl – pbs Tcl scheduler                                                                                                                                                                                                                                                  | 9                                                                                                                                                                                                                                                                                                                                   |
| <b>SYNOPSIS</b>                                                                                                                                                                                                                                                                    | 9                                                                                                                                                                                                                                                                                                                                   |
| pbs_sched [-a alarm] [-b file] [-d home] [-i file] [-L logfile] [-p file] [-S port] [-t file] [-v]                                                                                                                                                                                 | 9                                                                                                                                                                                                                                                                                                                                   |
| [-c file]                                                                                                                                                                                                                                                                          | 9                                                                                                                                                                                                                                                                                                                                   |
| <b>DESCRIPTION</b>                                                                                                                                                                                                                                                                 | 9                                                                                                                                                                                                                                                                                                                                   |
| The <b>pbs_sched</b> program runs in conjunction with the PBS server. It queries the server about the state of PBS and communicates with <b>pbs_mom</b> to get information about the status of running jobs, memory available etc. It then makes decisions as to what jobs to run. | 9                                                                                                                                                                                                                                                                                                                                   |
| pbs_sched must be executed with root permission.                                                                                                                                                                                                                                   | 9                                                                                                                                                                                                                                                                                                                                   |
| <b>OPTIONS</b>                                                                                                                                                                                                                                                                     | 9                                                                                                                                                                                                                                                                                                                                   |
| -a alarm                                                                                                                                                                                                                                                                           | This specifies the time in seconds to wait for a schedule run to finish. If a script takes too long to finish, an alarm signal is sent, and the scheduler is restarted. If a core file does not exist in the current directory, <b>abort()</b> is called and a core file is generated. The default for <i>alarm</i> is 180 seconds. |
| -b file                                                                                                                                                                                                                                                                            | This specifies the "body" file. The file given is read into memory once at program start or after the program receives a SIGHUP and executed each time the scheduler is awakened by the server. If this option is not given, the file "sched_tcl" in the directory PBS_HOME/sched_priv is read for the body code.                   |
| -d home                                                                                                                                                                                                                                                                            | This specifies the PBS home directory, PBS_HOME. The current working directory of the scheduler is PBS_HOME/sched_priv. If this option is not given, PBS_HOME defaults to \$PBS_SERVER_HOME as defined during the PBS build procedure.                                                                                              |
| -i file                                                                                                                                                                                                                                                                            | This specifies the "initialize" file. The file given is executed once before the main processing loop is entered. If this option is not given, no initialization code is executed.                                                                                                                                                  |
| -L logfile                                                                                                                                                                                                                                                                         | Specifies an absolute path name of the file to use as the log file. If not specified, the scheduler will open a file named for the current date in the PBS_HOME/sched_logs directory (see the -d option).                                                                                                                           |
| -p file                                                                                                                                                                                                                                                                            | This specifies the "print" file. Any output from the Tcl code which is written to standard out or standard error will be written to this file. If this option is not given, the file used will be <i>PBS_HOME/sched_priv/sched_out</i> . See the -d option.                                                                         |
| -S port                                                                                                                                                                                                                                                                            | This specifies the port to use. If this option is not given, the default port for the PBS scheduler is used.                                                                                                                                                                                                                        |
| -t file                                                                                                                                                                                                                                                                            | This specifies the "terminator" file. If a QUIT command is sent from the server, this code is executed before the scheduler exits. If this option is not given, no special termination handling is done.                                                                                                                            |
| -v                                                                                                                                                                                                                                                                                 | This puts the scheduler into "verbose" mode. Any errors will be shown no matter what this may be set to, but some "uninteresting" events may be logged by using this flag. An example is a message each time the server contacts the scheduler.                                                                                     |
| -c file                                                                                                                                                                                                                                                                            | Specify a configuration file, see description below. If this is a relative file name it will be relative to PBS_HOME/sched_priv, see the -d option. If the -c option is not supplied, pbs_sched will not attempt to open a configuration file.                                                                                      |

The options that specify file names may be absolute or relative. If they are relative, their root directory will be PBS\_HOME/sched\_priv. 9  
9

## USAGE 9

This version of the scheduler requires knowledge of the Tcl language. 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. 9  
9  
9  
9  
9  
9

openrm host ?port? 9

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. 9  
9  
9

closerm connection 9

The parameter *connection* is a handle to a resource monitor which was previously returned from **openrm**. This connection is closed. Nothing is returned. 9  
9

downrm connection 9

Sends a command to the connected resource monitor to shutdown. Nothing is returned. 9  
9

configrm connection filename 9

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. 9  
9

addreq connection request 9

A resource request is sent to the connected resource monitor. If this is successful, a "0" is returned, otherwise, "-1" is returned. 9  
9

getreq connection 9

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. 9  
9

allreq request 9

A resource request is sent to all connected resource monitors. The number of streams acted upon is returned. 9  
9

flushreq 9

All resource requests previously sent to all connected resource monitors are flushed out to the network. Nothing is returned. 9  
9

activereq 9

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. 9  
9

fullresp flag 9

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. 9  
9  
9  
9  
9

pbsstatserv 9

The server is sent a status request for information about the server itself. If the re- 9

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |                                      |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------|
| quest 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).                                                                                                                                                                                                                                                                                                  | 9<br>9<br>9                          |
| <b>pbsstatjob</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | 9                                    |
| 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). | 9<br>9<br>9<br>9<br>9<br>9<br>9<br>9 |
| <b>pbsstatque</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | 9                                    |
| 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 <b>pbsstatjob</b> . The third is the "text" associated with the queue (usually blank).                                                                                                | 9<br>9<br>9<br>9<br>9<br>9           |
| <b>pbsstatnode</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | 9                                    |
| 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 nodes's name. The second is a list of attributes similar to <b>pbsstatjob</b> . The third is the "text" associated with the node (usually blank).                                                                                                    | 9<br>9<br>9<br>9<br>9<br>9           |
| <b>pbsselstat</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | 9                                    |
| 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 <b>pbsstatjob</b> is returned, otherwise an empty string is returned.                                                                                                                                                                                                                                                                                                                    | 9<br>9<br>9                          |
| <b>pbsrunjob jobid ?location?</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | 9                                    |
| Run the job given by <i>jobid</i> at the location given by <i>location</i> . If <i>location</i> is not given, the default location is used. If this is successful, a "0" is returned, otherwise, "-1" is returned.                                                                                                                                                                                                                                                                                                                              | 9<br>9<br>9                          |
| <b>pbsasyrunjob jobid ?location?</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | 9                                    |
| Run the job given by <i>jobid</i> at the location given by <i>location</i> without waiting for a positive response that the job has actually started. If <i>location</i> is not given, the default location is used. If this is successful, a "0" is returned, otherwise, "-1" is returned.                                                                                                                                                                                                                                                     | 9<br>9<br>9                          |
| <b>pbsrerunjob jobid</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | 9                                    |
| Re-runs the job given by <i>jobid</i> . If this is successful, a "0" is returned, otherwise, "-1" is returned.                                                                                                                                                                                                                                                                                                                                                                                                                                  | 9<br>9                               |
| <b>pbsdeljob jobid</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | 9                                    |
| Delete the job given by <i>jobid</i> . If this is successful, a "0" is returned, otherwise, "-1" is returned.                                                                                                                                                                                                                                                                                                                                                                                                                                   | 9<br>9                               |
| <b>pbsholdjob jobid</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | 9                                    |
| Place a hold on the job given by <i>jobid</i> . If this is successful, a "0" is returned, otherwise, "-1" is returned.                                                                                                                                                                                                                                                                                                                                                                                                                          | 9<br>9                               |
| <b>pbsmovejob jobid ?location?</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | 9                                    |
| Move the job given by <i>jobid</i> to the location given by <i>location</i> . If <i>location</i> is not given, the default location is used. If this is successful, a "0" is returned, otherwise, "-1" is returned.                                                                                                                                                                                                                                                                                                                             | 9<br>9<br>9                          |

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |   |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| <code>pbsqenable queue</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | 9 |
| Set the "enabled" attribute for the queue given by <i>queue</i> to true. If this is successful, a "0" is returned, otherwise, "-1" is returned.                                                                                                                                                                                                                                                                                                                                                                                                                             | 9 |
| <code>pbsqdisable queue</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | 9 |
| Set the "enabled" attribute for the queue given by <i>queue</i> to false. If this is successful, a "0" is returned, otherwise, "-1" is returned.                                                                                                                                                                                                                                                                                                                                                                                                                            | 9 |
| <code>pbsqstart queue</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | 9 |
| Set the "started" attribute for the queue given by <i>queue</i> to true. If this is successful, a "0" is returned, otherwise, "-1" is returned.                                                                                                                                                                                                                                                                                                                                                                                                                             | 9 |
| <code>pbsqstop queue</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | 9 |
| Set the "started" attribute for the queue given by <i>queue</i> to false. If this is successful, a "0" is returned, otherwise, "-1" is returned.                                                                                                                                                                                                                                                                                                                                                                                                                            | 9 |
| <code>pbsalterjob jobid attribute_list</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | 9 |
| Alter the attributes for a job specified by <i>jobid</i> . The parameter <i>attribute_list</i> 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.                                                                                                                                                                                         | 9 |
| <code>pbsrescquery resource_list</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | 9 |
| Obtain information about the resources specified by <i>resource_list</i> . This will be a list of strings. If the request succeeds, a list with the same number of elements as <i>resource_list</i> is returned. Each element in this list will be a list with four numbers. The numbers specify <i>available</i> , <i>allocated</i> , <i>reserved</i> , and <i>down</i> in that order.                                                                                                                                                                                     | 9 |
| <code>pbsrescreserve resource_id resource_list</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | 9 |
| Make (or extend) a reservation for the resources specified by <i>resource_list</i> which will be given as a list of strings. The parameter <i>resource_id</i> is a number which provides a unique identifier for a reservation being tracked by the server. If <i>resource_id</i> 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. | 9 |
| <code>pbsrescrelease resource_id</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | 9 |
| The reservation specified by <i>resource_id</i> is released.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | 9 |
| 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.                                                                                                                                                                                                                                                                                                         | 9 |
| <code>pbsconnect ?server?</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | 9 |
| 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.                                                                                                                                                                                                                                                                                                                                                                                                                        | 9 |
| <code>pbsdisconnect</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | 9 |
| Disconnect from the currently connected server.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | 9 |
| The above Tcl functions use PBS interface library calls for communication with the server and the PBS resource monitor library to communicate with <code>pbs_mom</code> .                                                                                                                                                                                                                                                                                                                                                                                                   | 9 |
| <code>datetime ?day? ?time?</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | 9 |
| 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.                                                                                                                                                                                                                                                                                                                                                                                                             | 9 |
| 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:                                                                                                                                                                                                                                                                                                                                                                                                                          | 9 |

YYMMDDhhmmss 9

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. 9

The second option with one argument is a relative time. The format for this is HH:MM:SS 9

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. 9

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. 9

**strftime format time** 9  
 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 second 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. 9

**logmsg tag message** 9  
 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. 9

The Tcl interpreter is started at program initialization and after a reset (the receipt of a SIGHUP signal). It is not deleted between scheduling runs so variables which are set in one can be accessed later. 9

The "initialize" and "terminator" files are run with no supplied connection to the server. This means that none of the above functions which talk to the server will work unless **pbsconnect** is called first. The "body" file is run with a connection to the server already established. 9

**CONFIGURATION FILE** 9

A configuration file may be specified with the -c option. This file may be used to specify the hosts (servers) which are allowed to connect to pbs\_sched. The hosts are specified in the configuration file in a manor identical to that used in pbs\_mom. There is one line per host with the syntax: 9

\$clienthost hostname 9

where clienthost and hostname are separated by white space. 9

Two host names are always allowed to connection to pbs\_sched, "localhost" and the name returned to pbs\_sched by the system call *gethostname()*. These names need not be specified in 9



|                                                                                                                                               |   |
|-----------------------------------------------------------------------------------------------------------------------------------------------|---|
| the configuration file.                                                                                                                       | 9 |
| The configuration file must be "secure". It must be owned by a user id and group id less than 10 and not be world writable.                   | 9 |
| <b>FILES</b>                                                                                                                                  | 9 |
| SPBS_SERVER_HOME/sched_priv                                                                                                                   | 9 |
| the default directory for configuration files, typically (/usr/spool/pbs)/sched_priv.                                                         | 9 |
| <b>Signal Handling</b>                                                                                                                        | 9 |
| A C based scheduler will handle the following signals:                                                                                        | 9 |
| <b>SIGHUP</b>                                                                                                                                 | 9 |
| The server will close and reopen its log file and reread the config file if one exists.                                                       | 9 |
| <b>SIGALRM</b>                                                                                                                                | 9 |
| If the site supplied scheduling module exceeds the time limit, the Alarm will cause the scheduler to attempt to core dump and restart itself. | 9 |
| <b>SIGINT and SIGTERM</b>                                                                                                                     | 9 |
| Will result in an orderly shutdown of the scheduler.                                                                                          | 9 |
| All other signals have the default action installed.                                                                                          | 9 |
| <b>EXIT STATUS</b>                                                                                                                            | 9 |
| Upon normal termination, an exit status of zero is returned.                                                                                  | 9 |
| <b>SEE ALSO</b>                                                                                                                               | 9 |
| pbs_scheduler_cc(8B), pbs_scheduler_rule(8B), pbs_server(8B), pbs_mom(8B)                                                                     | 9 |

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |   |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| <b>7.2. Node Management</b>                                                                                                                                                                                                                                                                                                                                                                                                                                              | 9 |
| <b>NAME</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                              | 9 |
| pbsnodes – pbs node manipulation                                                                                                                                                                                                                                                                                                                                                                                                                                         | 9 |
| <b>SYNOPSIS</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                          | 9 |
| pbsnodes [-{c d o r}] [-s server] nodename ...                                                                                                                                                                                                                                                                                                                                                                                                                           | 9 |
| pbsnodes -{a d l} [-s server]                                                                                                                                                                                                                                                                                                                                                                                                                                            | 9 |
| <b>DESCRIPTION</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                       | 9 |
| The <b>pbsnodes</b> command is used to query the status of nodes, or mark nodes down, free or off-line. Node information is obtained by sending a request to the PBS job server.                                                                                                                                                                                                                                                                                         | 9 |
| If <b>pbsnodes</b> is run without options (other than <b>-s</b> ), and with a list of node names, then the status of the specified nodes is returned.                                                                                                                                                                                                                                                                                                                    | 9 |
| If pbsnodes is run without options (other than -s), and without a node list, the command usage syntax is printed.                                                                                                                                                                                                                                                                                                                                                        | 9 |
| Executing pbsnodes with the <b>-c</b> , <b>-d</b> , <b>-o</b> , or <b>-r</b> options requires the user to have PBS Manager or Operator privilege.                                                                                                                                                                                                                                                                                                                        | 9 |
| <b>OPTIONS</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                           | 9 |
| -a All nodes and all their attributes are listed.                                                                                                                                                                                                                                                                                                                                                                                                                        | 9 |
| -c Clear OFFLINE or DOWN from listed nodes. The listed nodes are "free" to be allocated to jobs.                                                                                                                                                                                                                                                                                                                                                                         | 9 |
| -d The nodes specified as operands are marked DOWN and unavailable to run jobs. It is important that all the nodes known to be down are given as arguments on the command line. This is because nodes which are not listed are assumed to be UP and will be indicated as such if they were previously marked DOWN. I.e., "pbsnodes -d" will mark all nodes as free. None of the other options will change the marking for nodes which are not given on the command line. | 9 |
| -l List all nodes marked in any way.                                                                                                                                                                                                                                                                                                                                                                                                                                     | 9 |
| -o Mark listed nodes as OFFLINE even if currently in use. This is different from being marked DOWN. An automated script that checks nodes being up or down and calls <b>pbsnodes</b> with a list of nodes down will not change the status of nodes marked OFFLINE. This gives the administrator a tool to hold a node out of service without changing the automatic script.                                                                                              | 9 |
| -r clear OFFLINE from listed nodes.                                                                                                                                                                                                                                                                                                                                                                                                                                      | 9 |
| -s specify the PBS server to which to connect.                                                                                                                                                                                                                                                                                                                                                                                                                           | 9 |
| <b>SEE ALSO</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                          | 9 |
| pbs_server(8B) and qmgr(8B)                                                                                                                                                                                                                                                                                                                                                                                                                                              | 9 |

|                                                                                                                                                                                                                                  |   |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| <b>7.3. Manage Server</b>                                                                                                                                                                                                        | 9 |
| <b>NAME</b>                                                                                                                                                                                                                      | 9 |
| qmgr – pbs batch system manager                                                                                                                                                                                                  | 9 |
| <b>SYNOPSIS</b>                                                                                                                                                                                                                  | 9 |
| qmgr [-a] [-c command] [-e] [-n] [-z] [server...]                                                                                                                                                                                | 9 |
| <b>DESCRIPTION</b>                                                                                                                                                                                                               | 9 |
| The <b>qmgr</b> command provides an administrator interface to the batch system.                                                                                                                                                 | 9 |
| The command reads directives from standard input. The syntax of each directive is checked and the appropriate request is sent to the batch server or servers.                                                                    | 9 |
| The list or print subcommands of qmgr can be executed by general users. Creating or deleting a queue requires PBS Manager privilege. Setting or unsetting server or queue attributes requires PBS Operator or Manager privilege. | 9 |
| <b>OPTIONS</b>                                                                                                                                                                                                                   | 9 |
| -a Abort <b>qmgr</b> on any syntax errors or any requests rejected by a server.                                                                                                                                                  | 9 |
| -c <i>command</i> Execute a single <i>command</i> and exit <b>qmgr</b> . The command must be enclosed in quote marks, e.g.<br>qmgr -c "print server"                                                                             | 9 |
| -e Echo all commands to standard output.                                                                                                                                                                                         | 9 |
| -n No commands are executed, syntax checking only is performed.                                                                                                                                                                  | 9 |
| -z No errors are written to standard error.                                                                                                                                                                                      | 9 |
| <b>OPERANDS</b>                                                                                                                                                                                                                  | 9 |
| The <i>server</i> operands identify the name of the batch server to which the administrator requests are sent. Each <i>server</i> conforms to the following syntax:<br>host_name[:port]                                          | 9 |
| where <i>host_name</i> is the network name of the host on which the server is running and <i>port</i> is the port number to which to connect. If <i>port</i> is not specified, the default port number is used.                  | 9 |
| If <i>server</i> is not specified, the administrator requests are sent to the local server.                                                                                                                                      | 9 |
| <b>STANDARD INPUT</b>                                                                                                                                                                                                            | 9 |
| The <b>qmgr</b> command reads standard input for directives until end of file is reached, or the <i>exit</i> or <i>quit</i> directive is read.                                                                                   | 9 |
| <b>STANDARD OUTPUT</b>                                                                                                                                                                                                           | 9 |
| If Standard Output is connected to a terminal, a command prompt will be written to standard output when qmgr is ready to read a directive.                                                                                       | 9 |
| If the <i>-e</i> option is specified, <b>qmgr</b> will echo the directives read from standard input to standard output.                                                                                                          | 9 |
| <b>STANDARD ERROR</b>                                                                                                                                                                                                            | 9 |
| If the <i>-z</i> option is not specified, the qmgr command will write a diagnostic message to standard error for each error occurrence.                                                                                          | 9 |
| <b>EXTENDED DESCRIPTION</b>                                                                                                                                                                                                      | 9 |
| If <b>qmgr</b> is invoked without the <i>-c</i> option and standard output is connected to a terminal, qmgr will write a prompt to standard output and read a directive from standard input.                                     | 9 |

Commands can be abbreviated to their minimum unambiguous form. A command is terminated by a new line character or a semicolon, ";", character. Multiple commands may be entered on a single line. A command may extend across lines by escaping the new line character with a back-slash "\".

Comments begin with the # character and continue to end of the line. Comments and blank lines are ignored by qmgr.

**DIRECTIVE SYNTAX** 9

A qmgr directive is one of the following forms: 9

```
command server [names] [attr OP value[,attr OP value,...]] 9
command queue [names] [attr OP value[,attr OP value,...]] 9
command node [names] [attr OP value[,attr OP value,...]] 9
```

Where, 9

command is the command to perform on a object. Commands are: 9

- active sets the active objects. If the active objects are specified, and the name is not given in a qmgr cmd the active object names will be used. 9
- create is to create a new object, applies to queues and nodes. 9
- delete is to destroy an existing object, applies to queues and nodes. 9
- set is to define or alter attribute values of the object. 9
- unset is to clear the value of attributes of the object. Note, this form does not accept an OP and value, only the attribute name. 9
- list is to list the current attributes and associated values of the object. 9
- print is to print all the queue and server attributes in a format that will be usable as input to the qmgr command. 9

names is a list of one or more names of specific objects The name list is in the form: 9  
 [name][@server][,queue\_name[@server]...] 9  
 with no intervening white space. The name of an object is declared when the object is first created. If the name is @server, then all the objects of specified type at the server will be effected. 9

attr specifies the name of an attribute of the object which is to be set or modified. The attributes of objects are described in section 2 of the ERS. If the attribute is one which consist of a set of resources, then the attribute is specified in the form: 9  
 attribute\_name.resource\_name 9

OP operation to be performed with the attribute and its value: 9  
 = set the value of the attribute. If the attribute has a existing value, the current value is replaced with the new value. 9  
 += increase the current value of the attribute by the amount in the new value. 9  
 -= decrease the current value of the attribute by the amount in the new value. 9

value the value to assign to an attribute. If the value includes white space, commas or other special characters, such as the # character, the value string must be inclosed in quote marks ("). 9

The following are examples of qmgr directives: 9

```
create queue fast priority=10,queue_type=e,enabled = true,max_running=0 9
set queue fast max_running +=2 9
create queue little 9
set queue little resources_max.mem=8mw,resources_max.cput=10 9
```

```
unset queue fast max_running 9
set node state = "down,offline" 9
active server s1,s2,s3 9
list queue @server1 9
set queue max_running = 10 - uses active queues 9
```

**EXIT STATUS** 9

Upon successful processing of all the operands presented to the qmgr command, the exit status will be a value of zero. 9

If the qmgr command fails to process any operand, the command exits with a value greater than zero. 9

**SEE ALSO** 9

pbs\_server(8B), pbs\_queue\_attributes(7B), pbs\_server\_attributes(7B), qstart(8B), qstop(8B), qenable(8B), and qdisable(8). 9

**8. The PBS Job Scheduler** 9

The PBS Job Scheduler is a daemon that performs job scheduling for a particular PBS server. 9

**8.1. Scheduler Overview** 9

The PBS Scheduler performs job scheduling upon request from a PBS Server. Except for the initial message, which is transmitted through a normal socket connection, the Server and Scheduler communicate by using functions defined in the `pbs_ifl` (PBS interface library). 9

The entity which is given the role of scheduler for PBS may be a special purpose program or even a shell script. The fundamental parts have been provided to offer maximum flexibility. Two example schedulers have been included. One scheduler is based on **yacc** and **lex** and creates its own language BASL, for BAtch Scheduler Language. The other scheduler is based on **tcl** and augments an already existing language with some new commands that give access to PBS functions. The other is a C language framework. 9

**8.2. The Scheduling Cycle** 9

If the scheduler has been asked to initiate a scheduling cycle, it first reverses the sense of the connection between the scheduler and server, and requests information concerning the current job pool from the server, and possibly, information about the queues that the server has defined. The scheduler then evaluates the job set according to the procedure expressed in the language file, and determines if any action should be taken concerning any jobs in the pool. At the end of the cycle, the scheduler once more resumes listening on the designated port. 9

The scheduler may also request information from one or more Resource Monitors. Resource Monitors are daemons that provide information about the state of the resources of particular machines, or collections of processors, such as clusters or MPP partitions. Communication between the scheduler and resource Monitors is conducted over a well-known port. The protocol used between the scheduler and a resource Monitor is described in the PBS Scheduler IDS. 9

**8.3. The Tcl Scheduler** 9

Tcl is a flexible language which can be extended fairly easily. The library to communicate with the Resource Monitor and the regular PBS Interface Library have been wrapped into Tcl commands to create a special purpose tcl interpreter which waits to be awakened by the PBS server before swinging into action. The indispensable book for dealing with Tcl is *Tcl and the Tk Toolkit* by John K. Ousterhout. 9

The Tcl scheduler man page gives all the particulars about the functions which are available to access the PBS Resource Monitor and Server. The section of the ERS which describes the Resource Monitor discusses the information available to the script writer for a particular type of system. 9

**8.4. The C Scheduler** 9

This is a skeleton program designed as a starting point for someone to use to create a complete C language scheduler program. It consists of a main program that processes options and sets up the network connection to the server. It was patterned after the Tcl scheduler main program in the hope it would not need to be changed yet the scheduler algorithm could still be modified easily. A directory of examples is included that show the required entry points for the scheduler writer. 9

[Page intentionally left blank.]

## 9. Resource Monitor

This section describes the PBS resource monitor functions of the Machine Oriented Miniserver (MOM). MOM is the part of PBS which is machine dependent rather than pure POSIX. The resource monitor communicates with the world (mainly the PBS scheduler) over the network. All the resource requests discussed below should be sent in as large a group as practical. This helps reduce the impact on the machine answering the request because `pbs_mom` tries to gather information only once for each request. Each resource reported within each request will then also be consistent with all others.

Each resource request to `pbs_mom` is sent with a command number telling what the rest of the request will contain. For the command "tell me about these resources", the following strings are each a separate resource request. These consist of a resource name with parameters inclosed in square brackets. Each parameter has a parameter name followed by an equal sign and then a value. For example `size[fs=/tmp]`.

### 9.1. Resources

The machines and operating systems supported by `pbs_mom` are

- Cray systems using *Unicos 8, 9 or 10, or Unicos MK 2*
- IBM 590 Workstations using *AIX 4*
- IBM SP using *AIX 4* with *PSSP 2.1* or *PSSP 3.1*.
- Silicon Graphics systems using *IRIX 5.x* or *6.x*,
- Sun Sparc using *SunOS 4.1.x* or *Solaris 2.5 (5.5)*,
- Fujitsu VPP300 systems using *UXP/V 4.1 ES*,
- AMD/Intel/Cyrix systems using *Linux* or *FreeBSD*

All of these machines will support a standard list of resources.

`arch` This returns a string which specifies the machine architecture of the host being queried. By default, it will be one of: "aix4", "irix5", "fujitsu", "irix6", "irix6array", "linux", "bsd", "solaris5", "sunos4", "unicos8" or "unicosmk2". The string returned may be changed by adding a *static resource* or *shell command* to Mom's configuration file. The value of the configuration file entry will override that of the standard return for `arch`.

Note, an configure file entry with the same name as any standard resource will override the standard value when reported to a resource request from a privileged port on a *client host*. A resource request from an unprivileged port on a *restricted host* will receive the standard value; the configure file entries are ignored.

`cput` This reports cpu time in seconds. Two different parameters are accepted. One is *proc* which can be used to specify a process, the other is *session* which can be used to specify a session. For example: `cput[proc=8765]` and `cput[session=4567]`.

`idletime`  
This is the time in seconds in which no keystroke or mouse movement has taken place on any terminal connected to the system.

`loadave`  
This reports the smoothed system loadave (number of processes in the kernel run queue).

`mem` This reports memory usage in bytes. The *proc* and *session* parameters are accepted. For example: `mem[proc=8765]` and `mem[session=4567]`.



|                                                                                                                                                                                                         |                                                                                                                                                                                                                                                                                                      |   |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| <code>ncpus</code>                                                                                                                                                                                      | Returns the number of processors that are available.                                                                                                                                                                                                                                                 | 9 |
| <code>nsessions</code>                                                                                                                                                                                  | Returns the number of sessions which exist in the system. Sessions owned by root are excluded. No parameters are allowed.                                                                                                                                                                            | 9 |
| <code>nusers</code>                                                                                                                                                                                     | Reports the number of users who have processes running in the system, excluding root. No parameters are allowed.                                                                                                                                                                                     | 9 |
| <code>pids</code>                                                                                                                                                                                       | Reports a list of processes for a session. A parameter <i>session</i> must be given. For example <code>pids[session=34615]</code> .                                                                                                                                                                  | 9 |
| <code>physmem</code>                                                                                                                                                                                    | Returns the physical (main) memory size in kilobytes.                                                                                                                                                                                                                                                | 9 |
| <code>sessions</code>                                                                                                                                                                                   | This will return a list of the sessions which exist in the system. This is different from PBS "jobs" and the sessions listed may not be part of any PBS job. Sessions owned by root are excluded. No parameters are allowed. (Was "jobs" in prior versions)                                          | 9 |
| <code>size</code>                                                                                                                                                                                       | Reports the size of file system objects in kilobytes. Two different parameters may be given. One is <i>file</i> which specifies a filename whose size is returned. The other is <i>fs</i> which specifies a directory in a file system which is examined to find the amount of free space available. | 9 |
| <code>uname</code>                                                                                                                                                                                      | This returns a string with the POSIX specified information returned from the <code>uname()</code> function separated by spaces. The strings are in order: <i>sysname</i> , <i>nodename</i> , <i>release</i> , <i>version</i> and <i>machine</i> .                                                    | 9 |
| <code>validuser</code>                                                                                                                                                                                  | This function returns a string of either <i>yes</i> or <i>no</i> if the user name is valid, i.e. has a password entry. A parameter of <code>user=name</code> must be given. For example: <code>validuser[user=joe]</code> .                                                                          | 9 |
| <code>walltime</code>                                                                                                                                                                                   | Reports the time in seconds which a <i>proc</i> or <i>session</i> has existed in the system.                                                                                                                                                                                                         | 9 |
| The following resources formerly were available under prior versions of PBS built with {NEEDNODES} from <code>pbs_mom</code> : <b>avail</b> , <b>reserve</b> , <b>totpool</b> , and <b>usepool</b> .    |                                                                                                                                                                                                                                                                                                      | 9 |
| They are now available from the job server, <code>pbs_server</code> , via the standard PBS API. See the man page for <code>pbs_resquery()</code> for more information.                                  |                                                                                                                                                                                                                                                                                                      | 9 |
| The following are general examples of resource queries to <code>pbs_mom</code> :                                                                                                                        |                                                                                                                                                                                                                                                                                                      | 9 |
| <code>pids[session=456]</code>                                                                                                                                                                          |                                                                                                                                                                                                                                                                                                      | 9 |
| <code>mem[proc=123]</code>                                                                                                                                                                              |                                                                                                                                                                                                                                                                                                      | 9 |
| <code>size[file=/etc/passwd]</code>                                                                                                                                                                     |                                                                                                                                                                                                                                                                                                      | 9 |
| <code>idletime</code>                                                                                                                                                                                   |                                                                                                                                                                                                                                                                                                      | 9 |
| In the given order, these resource requests would return a list of pid's for the session "456", the memory usage for process "123", the size of the file "/etc/passwd", and the idletime of the system. |                                                                                                                                                                                                                                                                                                      | 9 |
| 9.1.1. SunOS Resources                                                                                                                                                                                  |                                                                                                                                                                                                                                                                                                      | 9 |
| On a Sun system running SunOS, PBS supports several resources in addition to the above list.                                                                                                            |                                                                                                                                                                                                                                                                                                      | 9 |
| <code>availmem</code>                                                                                                                                                                                   | Returns the virtual memory available to be used.                                                                                                                                                                                                                                                     | 9 |

|                                                                                                                                                                                                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |   |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| quota                                                                                                                                                                                                      | Returns information about disk quotas. Several parameters are used. The first must be <i>type</i> which can have a value of <i>harddata</i> , <i>softdata</i> , <i>currdata</i> , <i>hardfile</i> , <i>softfile</i> , <i>currfile</i> , <i>timedata</i> or <i>timefile</i> . The second must be <i>dir</i> which is used to specify the directory of the file system for which quota information is desired. The last parameter must be <i>user</i> which is used to specify the user name or id number whose quota information is to be retrieved. The type "harddata" returns the hard limit for data storage in kilobytes. The type "softdata" returns the warning limit for data storage in kilobytes. The type "currdata" returns the current usage of data storage in kilobytes. The type "hardfile" returns the hard limit for the number of files. The type "softfile" returns the warning limit for the number of files. The type "currfile" returns the current number of files. The type "timedata" returns the number of seconds that a user has left in the grace period for excessive disk use, or zero if the grace period is not active. The type "timefile" returns the number of seconds that a user has left in the grace period for having an excessive number of files, or zero if the grace period is not active. | 9 |
| resi                                                                                                                                                                                                       | Returns the resident memory size in kilobytes for a <i>proc</i> or <i>session</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | 9 |
| totmem                                                                                                                                                                                                     | Returns the total virtual memory which exists in the system in kilobytes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | 9 |
| 9.1.2. Digital Unix                                                                                                                                                                                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | 9 |
| The resources for systems running Digital Unix include the standard set plus:                                                                                                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | 9 |
| platform                                                                                                                                                                                                   | Returns the cpu type as a string.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | 9 |
| 9.1.3. FreeBSD Resources                                                                                                                                                                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | 9 |
| The FreeBSD resources include the standard set plus all those supported by the Sun under SunOS except totmem and availmem.                                                                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | 9 |
| 9.1.4. SGI Resources                                                                                                                                                                                       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | 9 |
| The Silicon Graphics resources for either Irix 5 or 6 include the standard set plus all those supported by the Sun under SunOS. If the MOM is built for "irix6array", an additional resource is available: |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | 9 |
| availmask                                                                                                                                                                                                  | Return a MAXCNODES-bit string with a '1' in each position where there are two CPUs available for a node.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | 9 |
| 9.1.5. Solaris Resources                                                                                                                                                                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | 9 |
| The Sun Solaris resources include the standard set with the addition of the following.                                                                                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | 9 |
| platform                                                                                                                                                                                                   | Returns the cpu type as a string, for example SUNW,Ultra-1.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | 9 |
| 9.1.6. Fujitsu Resources                                                                                                                                                                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | 9 |
| The Fujitsu UXP/V resources include the standard set with the addition of the following.                                                                                                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | 9 |
| totmem                                                                                                                                                                                                     | Returns the total virtual memory which exists in the system in kilobytes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | 9 |
| availmem                                                                                                                                                                                                   | Returns the virtual memory available to be used.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | 9 |

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |   |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| 9.1.7. Cray Unicos Resources                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | 9 |
| Cray Unicos has some fundamental differences from both the Sun and SGI. It does not use virtual memory and it has a much more sophisticated quota system. Some of the same resource names are used with slightly different meanings. All of the standard resources are included as well as all the resources of the SGI except "walltime". Unicos uses a periodic data gathering routine to calculate swap rate and cpu idle values. The time period that is used for each calculation is SAMPLE_DELTA which is set to 10 seconds by default.                                                                                                                                                                             | 9 |
| availmem                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | 9 |
| Returns the memory which is available for use by programs.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | 9 |
| cpuidle                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | 9 |
| This is the percent of idle time that all the processors have experienced within the previous SAMPLE_DELTA seconds. This is a time filtered value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | 9 |
| cpuguest                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | 9 |
| This is the percent of time that all the processors have spent running a guest operating system within the previous SAMPLE_DELTA seconds. This is a time filtered value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | 9 |
| cpusysw                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | 9 |
| This is the percent of time that all the processors have spent in system wait within the previous SAMPLE_DELTA seconds. This is a time filtered value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | 9 |
| cpuunix                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | 9 |
| This is the percent of time that all the processors have spent running kernel code within the previous SAMPLE_DELTA seconds. This is a time filtered value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | 9 |
| cpuuser                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | 9 |
| This is the percent of time that all the processors have spent running user code within the previous SAMPLE_DELTA seconds. This is a time filtered value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | 9 |
| totmem                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | 9 |
| Returns the total memory in kilobytes minus what is taken by the kernel.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | 9 |
| swapavail                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | 9 |
| The number of characters of free space in the swap device(s).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | 9 |
| swpinrate                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | 9 |
| The swap in activity within the previous SAMPLE_DELTA seconds in characters per second. This is also a time filtered value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | 9 |
| swpoutrate                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | 9 |
| The swap out activity within the previous SAMPLE_DELTA seconds in characters per second. This is also a time filtered value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | 9 |
| swprate                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | 9 |
| The swap activity within the previous SAMPLE_DELTA seconds in characters per second. This includes both swap in and swap out transfers and is time filtered such that previous values of <i>swprate</i> are used to smooth the current calculated value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | 9 |
| swptotal                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | 9 |
| The total number of characters in the swap device(s).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | 9 |
| swapused                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | 9 |
| The number of characters of active data in the swap device(s).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | 9 |
| quota                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | 9 |
| The same quota types as the Sun and SGI are supported by the C90. The cray supports several others as well. The additional types only operate if the resource monitor is compiled with the symbol <i>SRFS</i> . They are "snap_avail", "ares_avail", "res_total", "soft_res", "delta" and "reserve". Additionally, if <i>SRFS</i> is defined, the "dir" attribute can specify one of several special directories by starting with a dollar sign (\$) character. The allowed special names are \$TMPDIR, \$BIGDIR, \$FASTDIR or \$WRKDIR. These names are read from the file <i>/etc/tmpdir.conf</i> so in theory they could be changed but in practice they have been hardwired into MOM and the Server as resources. The | 9 |

file just gives the administrator the ability to change the actual directory where space will be allocated for a user making a SRFS request. 9  
9

If the "type" attribute is one of the SRFS values, there can be no other parameters. If the "type" attribute is one of the standard values, there must be one more parameter. It can have a name of "user", "group" or "account" and a value of a name or id number. The standard types have the same meaning as the other machines. The meanings of the SRFS types are very UNICOS specific. They are *snap\_avail*, *ares\_avail*, *res\_total*, *soft\_res*, *delta*, and *reserve*. The type "snap\_avail" returns the amount of total space, in characters. The type "ares\_avail" returns *snap\_avail* less unused job reserved space. The type "res\_total" returns the total amount of space reserved for jobs. The type "soft\_res" returns "true" if soft reservation is allowed and "false" otherwise. The type "delta" returns the setting of over or under commitment. The type "reserve" returns the number of characters committed to "srfs\_assist" mode jobs. For more information, see the UNICOS SRFS documentation. 9  
9  
9  
9  
9  
9  
9  
9  
9  
9

srfs\_reserve 9  
This gives the ability to set the number of characters committed to "srfs\_assist" mode jobs. The number is not a delta value so if there is more than one job in "srfs\_assist" mode, their requests must be added together before making this call. 9  
9

9.1.8. Cray Unicos MK 2 Resources 9

Cray Unicos MK 2 has some of the same resources as Unicos. Quite a few are missing because Unicos MK 2 typically comes in a "binary" release that does not include header files that allow programs such as PBS's MOM to retrieve information about processes in a general way. The resources available are "cput", "totmem", "availmem", "ncpus" (the comment for this says "Number of started processors" but it does not show the number of PEs), "physmem", "size", "idletime" and "quota". 9  
9  
9  
9  
9

In addition to these, PBS under Unicos/MK2 support the following resources: 9

mppe\_app 9  
The number of application PEs currently configured. 9

mppe\_free 9  
The total number of application PEs currently unallocated. 9

mppe\_avail 9  
The size of largest contiguous block of application PEs. 9

mppe\_info 9  
Returns the above three resources at the same time (intended to provide a scheduler interface with lower overhead). 9

9.1.9. IBM SP 9

Basically, the SP is a special case of AIX. However, because the SP Parallel Operating Environment works outside of PBS, PBS does not have access to memory usage or cpu time of tasks other than serial processes run from the job script. 9  
9

9.2. Libraries 9

A number of libraries exist to communicate with MOM. One is specialized as a Resource Monitor interface and another as a Task Manager interface. There is also a general network interface that uses UDP. These libraries are described below. 9  
9

### 9.2.1. Resource Monitor Library 9

The resource monitor library contains functions to facilitate communication with the resource monitor. It is set up to make it easy to connect to several resource monitors and handle the network communication efficiently. See the IDS for details. 9

### 9.2.2. Task Management Library 9

#### NAME 9

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 – task management API 9

#### SYNOPSIS 9

**#include <tm.h>** 9

**int tm\_init(info, roots)** 9

**void \*info;** 9

**struct tm\_roots \*roots;** 9

**int tm\_nodeinfo(list, nnodes)** 9

**tm\_node\_id \*\*list;** 9

**int \*nnodes;** 9

**int tm\_poll(poll\_event, result\_event, wait, tm\_errno)** 9

**tm\_event\_t poll\_event;** 9

**tm\_event\_t \*result\_event;** 9

**int wait;** 9

**int \*tm\_errno;** 9

**int tm\_notify(tm\_signal)** 9

**int tm\_signal;** 9

**int tm\_spawn(argc, argv, envp, where, tid, event)** 9

**int argc;** 9

**char \*\*argv;** 9

**char \*\*envp;** 9

**tm\_node\_id where;** 9

**tm\_task\_id \*tid;** 9

**tm\_event\_t \*event;** 9

**int tm\_kill(tid, sig, event)** 9

**tm\_task\_id tid;** 9

**int sig;** 9

**tm\_event\_t <sup>2</sup>event;** 9

**int tm\_obit(tid, obitval, event)** 9

**tm\_task\_id tid;** 9

**int \*obitval;** 9

**tm\_event\_t <sup>3</sup>event;** 9

**int tm\_taskinfo(node, tid\_list, list\_size, ntasks, event)** 9

**tm\_node\_id node;** 9

**tm\_task\_id \*tid\_list;** 9

**int list\_size;** 9

**int \*ntasks;** 9

**tm\_event\_t <sup>4</sup>event;** 9

**int tm\_atnode(tid, node)** 9

**tm\_task\_id tid;** 9

**tm\_node\_id \*node;** 9

```

int tm_rescinfo(node, resource, len, event)          9
    tm_node_id node;                                9
    char *resource;                                9
    int len;  9
    tm_event_t 5event;                                9
int tm_publish(name, info, len, event)            9
    char *name;                                    9
    void *info;                                    9
    int len;  9
    tm_event_t 6event;                                9
int tm_subscribe(tid, name, info, len, info_len, event) 9
    tm_task_id tid;                                  9
    char *name;                                    9
    void *info;                                    9
    int len;  9
    int *info_len;                                9
    tm_event_t 7event;                                9
int tm_finalize()                                  9

```

**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 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 implementaion 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 */  9
#define TM_ERROR_NODE ((tm_node_id)-1)                9

typedef int      tm_event_t; /* event handle, > 0 for real events */  9
#define TM_NULL_EVENT ((tm_event_t)0)                 9
#define TM_ERROR_EVENT((tm_event_t)-1)                9

```

```

typedef unsigned long tm_task_id;                      9
#define TM_NULL_TASK (tm_task_id)0                    9

```

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.

**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 {                                     9
    tm_task_id tm_me;                                 9
    tm_task_id tm_parent;                             9
    int tm_nnodes;                                   9
}

```

```

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

```

**tm\_me**                   The task id of this calling task. 9

**tm\_parent**             The task id of the task which spawned this task or TM\_NULL\_TASK if 9  
the calling task is the initial task started by PBS. 9

**tm\_nnodes**             The number of nodes allocated to the job. 9

**tm\_ntasks**             This will always be 0 for PBS. 9

**tm\_taskpoolid**         PBS does not support task pools so this will always be -1. 9

**tm\_tasklist**           This will be NULL for PBS. 9

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** be TM\_SUCCESS if the library initialization was successful, or an error return otherwise. 9

**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. 9

**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\_errno* will be set to an error code. 9

**tm\_notify()** is described in the PSCHED documentation, but is not implemented for PBS yet. It will return TM\_ENOTIMPLEMENTED. 9

**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. 9

**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*. 9

**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. 9

**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\_rescinfo()** 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 followed by a colon (:). 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. Each component has the resource name, an equal sign, and the limit value.

For example, a return for a task running on an SGI workstation might look like:

IRIX golum 6.2 03131015 IP22:cput=20:00,mem=400kb

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

**SEE ALSO**  
**pbs\_mom, PSCHED:**An <http://parallel.nas.nasa.gov/Psched/psched-api-report.ps>

### 9.2.3. Reliable Packet Protocol Library

**NAME**  
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 sockadd\_in \*addr;  
int rpp\_bind(port)  
int port;



|                                               |   |
|-----------------------------------------------|---|
| <b>int rpp_poll()</b>                         | 9 |
| <b>int rpp_io()</b>                           | 9 |
| <b>int rpp_read(stream, buf, len)</b>         | 9 |
| <b>u_int stream;</b>                          | 9 |
| <b>char *buf;</b>                             | 9 |
| <b>int len;</b>                               | 9 |
| <b>int rpp_write(stream, buf, len)</b>        | 9 |
| <b>u_int stream;</b>                          | 9 |
| <b>char *buf;</b>                             | 9 |
| <b>int len;</b>                               | 9 |
| <b>int rpp_close(stream)</b>                  | 9 |
| <b>u_int stream;</b>                          | 9 |
| <b>struct sockadd_in *rpp_getaddr(stream)</b> | 9 |
| <b>u_int stream;</b>                          | 9 |
| <b>int rpp_flush(stream)</b>                  | 9 |
| <b>u_int stream;</b>                          | 9 |
| <b>int rpp_terminate()</b>                    | 9 |
| <b>int rpp_shutdown()</b>                     | 9 |
| <b>int rpp_rcommit(stream, flag)</b>          | 9 |
| <b>u_int stream;</b>                          | 9 |
| <b>int flag;</b>                              | 9 |
| <b>int rpp_wcommit(stream, flag)</b>          | 9 |
| <b>u_int stream;</b>                          | 9 |
| <b>int flag;</b>                              | 9 |
| <b>int rpp_eom(stream)</b>                    | 9 |
| <b>u_int stream;</b>                          | 9 |
| <b>int rpp_getc(stream)</b>                   | 9 |
| <b>u_int stream;</b>                          | 9 |
| <b>int rpp_putc(stream, c)</b>                | 9 |
| <b>u_int stream;</b>                          | 9 |
| <b>int c;</b>                                 | 9 |

## 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)**.

**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. 9  
9

**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. 9  
9  
9  
9

**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. 9  
9  
9

**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 pointer is returned. 9  
9

**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. 9  
9

**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. 9  
9

**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. 9

**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. 9  
9  
9  
9  
9

**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. 9  
9  
9  
9

**rpp\_eom()** is called to terminate processing of the current message. 9

**SEE ALSO** 9  
**tcp(4P), udp(4P)** 9

[Page intentionally left blank.]

9

## 10. Security

9

Security in **PBS** is composed of two facets, authorization and authentication. Any user making a request of a batch server must be authorized to request that service. To provide the basis for authorization, each request contains the identity of the user making the request. To insure the user identity is correct, the identity must be authenticated in some manner.

### 10.1. Authorization

9

There is an authorization method, usually an *access control list*, or ACL, to control access to each object managed by a batch server, jobs, queues, and the server itself.

While the exact format of the entries in access control lists varies depending upon the list, they have several properties in common. Each list is maintained by the server as an editable text file. Each list entry consists of a single text line. Entries in the list may be prefixed with a single “+” or “-” character to indicate allowing or denying access. If neither character is present, allowing access is assumed. For example, a list containing:

```
-Kenneth
+Lora
-Julie
-Bob
```

Would explicitly grant access to Lora and explicitly deny access to Kenneth, Julie, and Bob.

Note that while each entry in the list is maintained in the file as a separate line, when entering the entries through **qmgr**, the entries are comma separated on a single line. For example:

```
set server some_list="-Kenneth,Lora,-Julie,-Bob"
```

The quote marks are required to correctly parse the string as a single value for the keyword “some\_list” as the string contains commas. Also be aware that when entered via qmgr, entries for ACLs which contain a host-domain part are sorted such that the more restrictive (fully specified, least wild carded) entries occur first. This is important to the comparison algorithm used when checking privilege. If the file is edited by hand, the order should be maintained.

#### 10.1.1. Server Access

9

Server access is controlled by a host access control list, and by two lists which grant special privilege to certain users.

##### 10.1.1.1. Server Host ACL

9

Each server has an access control list that specifies which hosts are allowed, or not allowed, to send requests to the server. The server depends on the network to identify the host which originated the connection (request). The entries in this list are fully specified host-domain names:

```
+host_one.domain.name.one
-host_two.domain.name.one
host_three.another.domain
*.my.domain
*.edu
```

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |   |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| - *                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | 9 |
| As shown in the last three entries, fields on the left, host, sub-domain, or domain names, may be “wild carded” by the use of a single asterisk. Thus any host in the domain <code>my.domain</code> or any host in any subdomain under the top level <code>edu</code> domain is granted access, while the last line denies access to any other host.                                                                                                                                                                  | 9 |
| <b>10.1.1.2. Managers and Operators</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | 9 |
| Each server has a list of “users at hosts” who have the privilege to request services restricted to operators and managers (administrators). Requests for those services from other users will be rejected. Requests from the named users on hosts other than the named host will also be rejected. The format of these lists is:                                                                                                                                                                                     | 9 |
| <code>user_name@host.domain.name</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | 9 |
| As above, the host, sub-domain, or domain name may be wild carded.                                                                                                                                                                                                                                                                                                                                                                                                                                                    | 9 |
| <b>10.1.2. Queue Access</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | 9 |
| Queue access is controlled by three levels: host, group, and user. Unlike other ACLs, the user and group list may have an extra first entry which determines the default for unnamed members. This entry in the list consists solely of a single “+” or “-” character with no additional characters. Any name, user or group, not specified will implicitly be granted or denied access based on this line. If multiple default lines, single + or -, appear in a list, it is undefined as to which takes precedence. | 9 |
| <b>10.1.2.1. Queue Host Access</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | 9 |
| As with the server, queue access can also be controlled by originating host. Each queue has an Host ACL identifying hosts from which jobs may or may not be entered into the particular queue. The format of the list entries is the same as the server’s host list.                                                                                                                                                                                                                                                  | 9 |
| <b>10.1.2.2. Queue Group Access</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | 9 |
| An access control list can be established which limits access to jobs run under certain group names. This is the group name under which the job would be executed. It may be different than the group name under which the user submitted the job, see the job attribute <code>group_list</code> . The format of this list is                                                                                                                                                                                         | 9 |
| <code>[+ -]</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | 9 |
| <code>[+ -]group_name</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | 9 |
| The optional first line of a single ‘+’ or ‘-’ establishes the default privilege for unlisted groups. If the first line is ‘-’, any unlisted group is denied access. If the first line is ‘+’ or contains additional characters, any unlisted group is allowed access.                                                                                                                                                                                                                                                | 9 |
| <b>10.1.2.3. Queue User Access</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | 9 |
| The user identification in each request may also control queue access. Each queue has an access control list identifying the users who are permitted or denied the ability to enter jobs into the queue. The format of this list is:                                                                                                                                                                                                                                                                                  | 9 |
| <code>[+ -]</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | 9 |
| <code>[+ -]user[@host.domain]</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | 9 |

The optional first line is identical to the group ACL. The “host.domain” may be wild carded as described under the Server Host ACL. If a user name is given without a “@host.domain”, then it is equivalent to that user name at any host, “user@\*”. 9  
9  
9

10.1.3. Job Access 9

The owner of a job is allowed full access to his or her own job. Operators and administrators also have access to all jobs. 9  
9

Users without special privilege, manager or operator, are not allowed to modify or delete a job that they do not own. The ability to query the status of a job of another user is determined by the server attribute query\_other\_jobs. If true, users may query the status of any user’s jobs. 9  
9  
9  
9

10.1.4. Root Jobs 9

Generally, the super user, root, is not allowed to submit batch jobs. This prevents someone who has cracked the system from doing additional damage. However, there is a build time option to allow root on the server’s local system to submit jobs. This option is controlled by the symbol PBS\_ROOT\_JOBS. Please see the server\_limits.h section of the PBS Administrator Guide for more information. 9  
9  
9  
9  
9

10.2. Authentication 9

There must be a guarantee that the identity of the user making a batch request is correct. A basic method described below is provided with **PBS**. It is called the PBS Interface Facility, or **pbs\_iff** (also known as Identify, Friend or Foe). The method of authenticating the user is somewhat modular. It can be removed and replaced by other authentication methods such as Kerberos with a bit of work in the server. 9  
9  
9  
9

In the basic method provided with **PBS**, an Authentication Batch Request is sent to the server from a privileged port. This request contains the user’s true identity, the name of the host on which the client is running, and the port name to which the client is bound and will be communicating with the server. This information is obtained by the *pbs\_iff* program. This program is a “setuid root” program allowing it to access privileged ports (those requiring root privilege) and thus giving some amount of credence to its claims. 9  
9  
9  
9  
9

The PBS user application programmable interface (API) library, **libpbs.a**, is the client side of the credential process. Whenever a user requests an open connection with a batch server through the API, the **pbs\_connect(3B)** routine will open a connection to the PBS server on a non-privileged port. It will then *fork()* and *exec()* **pbs\_iff** passing it the server’s name and the port to which *pbs\_connect()* is bound. The authentication type is returned by **pbs\_iff** to the *pbs\_connect* routine over a pipe. 9  
9  
9  
9

When the server receives the authentication request, it associates the user and host name with the network connection established by the client and identified by in the request. If the request is accepted by the server, the connection is considered authenticated for the remainder of its life. 9  
9  
9  
9

9



**11. Network Protocol** 9

The batch 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. 9

As new batch processing features will become part of the base, and as vendors will extend the product, the data exchange format must be extensible. 9

As the batch system will run on many systems within a complex, the data exchange format must provide for interoperability among systems with: 9

- Different architectures, including word size and byte order.
- Non uniform extensions to the batch systems.

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. 9

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. 9

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. 9

PBS originally used the ISO standard ASN.1 data encoding between all pieces of PBS except for the Scheduler to Resource Monitor communication. A public domain package, ISODE, provided for quick implementation and it was believed that ASN.1 would be required by POSIX when the batch protocol was standardized. In release 1.1.9, the user and administrative command utilities still used the ASN.1 encoding and the Job Server, *pbs\_server*, still accepted ASN.1 encoded requests from clients (user commands or other servers), but ASN.1 has being phased out in favor of the *Data Is Strings*, or DIS, encoding and is no longer used or supported. 9

**11.1. General DIS Data Encoding** 9

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. Human readability was a secondary goal and considered only when it would not cost. 9

Data are represented as people thing 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 or unsigned 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 9



even if it uses binary arithmetic. 9

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. 9

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. 9

The crux of the encoding is the integer. It's BNF is: 9

integer =: <count><sign><decimal string> | <count><integer> 9

count =: <decimal string> 9

sign =: '+' | '-' 9

decimal string =: <digit> | <digit><decimal string> 9

digit =: '0' | '1' | '2' | '3' | '4' | '5' | '6' | '7' | '8' | '9' 9

Single digit numbers do not require a count, just a sign followed by the one digit. Otherwise, the left most count is one digit long. Each count gives the number of digits in the decimal string to its right. The last decimal string is the actual integer value. The following table shows some examples of values and their encoding. 9

| Value      | Encoded String |
|------------|----------------|
| 1          | +1             |
| 2          | +2             |
| -3         | -3             |
| 12         | 2+12           |
| -13        | 2-13           |
| 1234567890 | 210+1234567890 |

The last value may take some explaining, the first digit, '2', is the first count and there for the length of the following integer. That integer, '10' is the second count. Since it is followed by a sign, its value follows and is therefore '10' digits in length. 9

Counted character strings (or simply strings) are defined as: 9

counted string =: <integer><characters> 9

characters =: <character> | <character><characters> 9

The number of characters in the "characters" value is given by "integer". Some examples: 9

| Value                    | Encoded String             |
|--------------------------|----------------------------|
| "abc"                    | +3abc                      |
| "This is a long string." | 2+22This is a long string. |

Real numbers are defined as: 9

real number =: <integer><integer> 9

The first integer is a real number is the coefficient, with no leading or trailing zeros and with the decimal point to its right. The second integer of a real number is the exponent to the base 10. 9

PBS provides a collection of the functions to perform the encoding/decoding of the basic data types. The collection is in the source tree under lib/Libdis. The objects are placed into the API library, libpbs.a, for ease of linking. This library may be used outside of PBS if desired. The routines are described (briefly) in the PBS Internal Design Specification. 9

Layered on top of the above routines, PBS makes use of a set of batch request and reply routines which use the DIS routines to encode/decode the data specific to the request or reply. This routines are called by the public PBS API routines defined in chapter 4 of this document. 9

**11.2. DIS Encoded Batch Requests** 9

The following describes the format of the PBS batch requests and replies using the DIS encoding. These requests are the basis of all communication between the client commands (qalter, qsub, ...) and the server, between two servers, and between the server and MOM. The Job Scheduler uses a different message format when talking with the Resource Manager side of MOM. 9

**11.2.1. Batch Requests** 9

Each request is encoded as a *Message Header*, a *Message Body*, and an *Message Extension*. The format of the Message Header and the Message Extension is the same for all batch requests. The Message Body varies for each request. 9

**11.2.1.1. Message Header** 9

unsigned integer: Batch Protocol Type Identifier (=2) 9  
 unsigned integer: Batch Protocol Version Identifier (=1) 9  
 unsigned integer: Batch Request Identifier 9  
 string: User Name 9

**11.2.1.2. Message Body** 9

Copy Files Request 9

string: Job Id 9  
 string: Job Owner (may be null string) 9  
 string: Execution User Name 9  
 string: Execution Group Name (may be null string) 9  
 unsigned integer: direction 9  
 unsigned integer: number of file pairs that follow 9  
 array of file pairs: 9  
     unsigned integer: file type flag 9  
     string: local (to mom) path name 9  
     string: remote path name 9

Job Credential Request 9

unsigned integer: credential type 9  
 string: opaque data 9

Job Files Move Request 9

unsigned integer: file block number 9  
 unsigned integer: file type (stdout, stderr, checkpoint, ...) 9  
 unsigned integer: amount of data in this block 9  
 string: Job Id 9  
 string: file data block 9

Job Obituary (Notice) Request 9

string: Job Id 9  
 unsigned integer: exit status 9  
 array of Attributes 9

|                                                                                                                                                                             |   |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| <b><u>Manager Request</u></b>                                                                                                                                               | 9 |
| <b>This request is used for a number of operations including Alter Job, Delete Job, Hold Job, and Release Job, as well as the request associated with the qmgr command.</b> | 9 |
| unsigned integer:      command (create, delete, set, unset, increment, decr)                                                                                                | 9 |
| unsigned integer:      object type                                                                                                                                          | 9 |
| string:                 object name                                                                                                                                         | 9 |
| array of Attributes                                                                                                                                                         | 9 |
| unsigned integer:      object type                                                                                                                                          | 9 |
| string:                 object name                                                                                                                                         | 9 |
| array of Attributes                                                                                                                                                         | 9 |
| <b><u>Message Job Request</u></b>                                                                                                                                           | 9 |
| string:                 Job Id                                                                                                                                              | 9 |
| unsigned integer:      which file (stderr, stdout, default)                                                                                                                 | 9 |
| string:                 message text                                                                                                                                        | 9 |
| <b><u>Move Job Request</u></b>                                                                                                                                              | 9 |
| string:                 Job Id                                                                                                                                              | 9 |
| string:                 Destination                                                                                                                                         | 9 |
| <b><u>Queue Job Request</u></b>                                                                                                                                             | 9 |
| string:                 Job Id                                                                                                                                              | 9 |
| string:                 Destination (queue name, null string implies default)                                                                                               | 9 |
| array of Attributes                                                                                                                                                         | 9 |
| <b><u>Register Dependency Request</u></b>                                                                                                                                   | 9 |
| string:                 Job Owner                                                                                                                                           | 9 |
| string:                 Parent Job Id                                                                                                                                       | 9 |
| string:                 Child Job Id                                                                                                                                        | 9 |
| unsigned integer:      dependency type                                                                                                                                      | 9 |
| unsigned integer:      operation                                                                                                                                            | 9 |
| signed long:           dependency cost                                                                                                                                      | 9 |
| <b><u>Run Job Request</u></b>                                                                                                                                               | 9 |
| string:                 Job Id                                                                                                                                              | 9 |
| string:                 Destination (where to run it)                                                                                                                       | 9 |
| <b><u>Shut Down Server Request</u></b>                                                                                                                                      | 9 |
| unsigned integer:      manner of shutdown                                                                                                                                   | 9 |
| <b><u>Signal Job Request</u></b>                                                                                                                                            | 9 |
| string:                 Job Id                                                                                                                                              | 9 |
| string:                 Signal name                                                                                                                                         | 9 |
| <b><u>Status Job, Queue, or Server Request</u></b>                                                                                                                          | 9 |

string: Object Id (or null string) 9  
 array of Attributes 9

Track Job (Notice) Request 9

string: Job Id 9  
 unsigned integer: hop count 9  
 string: location 9  
 unsigned character: current job state 9

**11.2.1.3. Complex Data Encoding** 9

Certain data items are complex, such as the array of attributes, and require special encoding as a set. 9

Array of Attributes – An array or list of attributes, regardless for job, queues, or a server, is encoded as follows: 9

unsigned integer: number of attributes in array (list) 9

followed by the attributes (if the above integer is non-zero): 9

unsigned integer: sum of length of the name, resource, and value strings, including the terminating null characters, 9  
 string: attribute name 9  
 unsigned integer: flag if resource name is present (1=yes, 0=no) 9  
 optional string: resource name (if above flag is 1) 9  
 string: value 9  
 unsigned integer: operation 9

Status Object (see Batch Reply) 9

Each status object in an array there of is encoded as: 9

unsigned integer: object type (server, queue, job) 9  
 string: object name 9  
 array of Attributes 9

**11.2.1.4. Message Extension** 9

unsigned integer: flag indicates presence of extension string (1=yes, 0=no) 9  
 optional string: extension string (present only if above flag is 1) 9

**11.2.2. Batch Reply Format** 9

The PBS batch reply is in two parts, a fixed header and a body which depends on the type of reply indicated in the header. 9

**11.2.2.1. Reply Header** 9

The reply header is encoded as: 9

unsigned integer: Batch Protocol type (2) 9  
 unsigned integer: Batch Protocol version (1) 9

|                   |                     |   |
|-------------------|---------------------|---|
| unsigned integer: | return code         | 9 |
| unsigned integer: | auxiliary code      | 9 |
| unsigned integer: | body type indicator | 9 |

#### 11.2.2.2. Reply Body 9

The reply bodies are encoded depending on type as: 9

Null body – no additional data 9

Job Id Reply – (used for Queue, Ready to Commit, and Commit requests) 9

|         |        |   |
|---------|--------|---|
| string: | Job Id | 9 |
|---------|--------|---|

Select Reply 9

|                   |                                   |   |
|-------------------|-----------------------------------|---|
| unsigned integer: | count of job id strings to follow | 9 |
| string:           | job id, repeated as required      | 9 |

Status Reply 9

|                         |                          |   |
|-------------------------|--------------------------|---|
| unsigned integer:       | count of status objects: | 9 |
| array of status objects |                          | 9 |

Text Reply 9

|         |      |   |
|---------|------|---|
| string: | text | 9 |
|---------|------|---|

Locate reply 9

|         |                 |   |
|---------|-----------------|---|
| string: | location of job | 9 |
|---------|-----------------|---|

### 11.3. Description of Data Exchange Format 9

A batch client makes a *service request* of a batch server. The request consists of one or more *subrequests* from the client to the server. For each subrequest, the server returns a *reply* to the client. As a minimum, the reply indicates success or failure of the subrequest. For some subrequests, such as a status request, additional information is included in the reply. 9

### 11.4. Request Format 9

A general request consists of two required parts, a *fixed header* and a *variable body*, and an optional third part, a *request extension* byte string. The header contains three (3) fields: 9

**request code** which is a number identifying the type of request. 9

**request user** a string which is the effective user name of the user making the request, 9

**request credential** an "any" structure which contains client authentication credentials. 9

The format and content of the request body depends on the type of request. The body is decoded separately from the header and after the header identifies the type of body. 9

The request extension byte string is intended to allow portable extensions to the batch command parameters. PBS makes use of this capability for the **qhold** and **qdel** commands. When a server receives a request extension that it does not recognize, it is free to abort the request.

**11.5. Reply Format** 9

The reply returned by a server to a client consists of a *fixed header* and an optional *variable body*. The body is only required in replies to certain types of requests.

The header contains the following integer numbers: 9

**reply code** which contains the primary success or failure code. 9

**auxiliary code** which contains extended status or failure information. 9

The optional body format depends on the type of request to which the reply is being returned. The format of the reply will be discussed under those requests which require the additional information contained in a body. 9

**11.6. Sequence of Subrequests** 9

The following sequence of requests is followed in the client - server communication. 9

**11.6.1. Open Connection** 9

The client opens a reliable stream connection with the server. The PBS mechanism for opening the connection is the **connect(2)** call to connect to an address associated with the requested server. 9

**11.6.2. Service Requests** 9

The client sends one or more service requests to the server. 9

The server returns a reply following processing of each request. The client does not send an additional request until it has received a reply to the prior request. 9

**11.6.3. Connection Closure** 9

The client closes the connection to the server to indicate that there are no more requests. The server may close the connection. If this occurs before the reply is sent, the client treats this as error condition. 9

**11.7. Description of Service Request** 9

The following sections describe the data items in each request. 9

**11.7.1. Queue Job Request** 9

**Queue Job** is a complex request, made up of four or more sub-requests: 9

- Queue Job (subrequest) 9
- Job Credentials 9
- Job Script, zero or more sub-requests. 9
- Ready to Commit. 9

|                                                                                                                                                                                                                                                                                                                                                                                                                                                        |   |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| – Commit ownership.                                                                                                                                                                                                                                                                                                                                                                                                                                    | 9 |
| The server must respond to each of the subrequests with a reply. The request may be aborted without the job being queued by the server at any point prior to receiving the Commit request.                                                                                                                                                                                                                                                             | 9 |
| 11.7.1.1. Queue Job                                                                                                                                                                                                                                                                                                                                                                                                                                    | 9 |
| The <b>Queue Job</b> subrequest informs the server that the client wishes the server to accept and enqueue a batch job.                                                                                                                                                                                                                                                                                                                                | 9 |
| The <i>DestinationID</i> field identifies the destination of the job. It is a string in the form: [queue][@server] If queue is not specified, the default queue is assumed. If @server is specified, it is ignored.                                                                                                                                                                                                                                    | 9 |
| If the optional field <i>JobIdentifier</i> is present, the job is an existing job being moved from one server to another. The <i>JobIdentifier</i> field contains the job identifier assigned to the job by the creating server. If <i>JobIdentifier</i> is not present, the job is being newly submitted by a client qsub. The receiving server will assign a job identifier to the job. The server, in either case, replies with the job identifier. | 9 |
| Also specified in the Queue Job subrequest are the batch job's private and public attributes, including required resources and the inherited environment variables, the <i>Variable_List</i> .                                                                                                                                                                                                                                                         | 9 |
| 11.7.1.2. Job Credential                                                                                                                                                                                                                                                                                                                                                                                                                               | 9 |
| The Job Credential is an optional subrequest. The request is provided as a means of passing service credentials or tickets for use by the batch job. These services are system or network services unrelated to the batch system. An example is a ticket authenticating the batch job to AFS (Andrew File System).                                                                                                                                     | 9 |
| If present, the server will just decode the data from the network presentation. Sites or vendors must add the additional functions to the server to present the credentials to the system when the job is placed into execution.                                                                                                                                                                                                                       | 9 |
| 11.7.1.3. Job Script                                                                                                                                                                                                                                                                                                                                                                                                                                   | 9 |
| The body of the <b>Job Script</b> request consists of sections of the job script file supplied to the qsub command. Each section contains up to 8192 bytes of the file. The last section may be shorter than 8192 bytes. The sections are sent in order and each must be acknowledged before the next is sent.                                                                                                                                         | 9 |
| Each Job Script sub-request also contains a sequence number. This integer starts with one (1) and is incremented by one for each additional Job Script sub-request.                                                                                                                                                                                                                                                                                    | 9 |
| 11.7.1.4. Ready to Commit                                                                                                                                                                                                                                                                                                                                                                                                                              | 9 |
| Upon receipt of the Ready to Commit, the receiving server <b>must</b> record the job to permanent storage <b>before</b> acknowledging.                                                                                                                                                                                                                                                                                                                 | 9 |
| The reply to the Ready to Commit, contains the job identifier. This is provided to assist in the recovery of a server or network failure.                                                                                                                                                                                                                                                                                                              | 9 |
| If the client is another batch server which had the batch job queued, upon receipt of a acknowledgment to the Ready to Commit, the client removes the job from its queue and permanent storage, deleting the job.                                                                                                                                                                                                                                      | 9 |
| 11.7.1.5. Commit                                                                                                                                                                                                                                                                                                                                                                                                                                       | 9 |

When the sending client has received a positive reply to the *Ready to Commit* request and had deleted the job locally, then and only then should the client send a *Commit* sub-request to the server. This ensures the job cannot be duplicated. 9

When the server receives the commit sub-request, the server assumes ownership of the job and perform the steps required to place the job into a queue. The server **does not reply** to the client following the commit sub-request. 9

11.7.1.6. Error Recovery 9

The following procedures are used to ensure the job is not lost or duplicated if the event of a communication failure between the sending client and the receiving server: 9

Client Side Recovery - Failure of Server 9

The follow procedures should be followed by the client attempting to send a job to a server when communications are lost with the server. 9

1. If the communication link is lost before the client sends the *Ready to Commit*, it is to assume the job is lost and restart from the beginning. The server will replace the job with the new copy when the client re-established the connection. 9
2. If the communication link is lost after the client sends the *Ready to Commit* but before sending the *Commit*, the client can assume that prior job information is at the server. Upon re-establishment of a connection with the server, the client re-sends the *Ready to Commit* followed by the *Commit*. 9
3. If the communication is lost after sending the *Ready to Commit* and at a later time the client receives a positive acknowledgment to the *Ready to Commit*, even if one has already been received, the client sends or resends the *Commit*. Note, the job identifier is required in the acknowledgment for this purpose. Also note that at this point the client should not have and does not need a copy of the job. 9
4. If the client is a temporary process, for example qsub, rather than restarting the Queue Job request, it may abort and issue an error message to the user. If the problem occurs after the *Ready to Commit* is sent but before an acknowledgment is received, it is unclear if the server has the complete job. The user should be warned to check for the job before attempting to resubmit it. Once the acknowledgment for the *Ready to Commit* has been received by the temporary client, the client may safely assume the job has been queued by the server. 9

Client Side Recovery - Failure of Client 9

Should the client side crash, upon recovery a permanent process (a sending server) performs the actions described in 1 through 3 above. 9

Server Side Recovery - Failure of Client 9

The following recovery procedures should be followed by a server receiving a job upon loss of communications with the client. 9

1. If the failure occurs before the *Ready to Commit* is received, the server discards the job. The client must restart from the beginning. 9
2. If the failure occurs after the *Ready to Commit* is received, the server should have recorded the job in permanent storage. The server keeps the job until (a) a request is received to delete it, or (b) the client resends the Read to Commit, Commit sequence. If after a "site defined" period of time, the server has not received any directions, it may notify the batch administrator and request instruction. 9

Server Side Recovery - Failure of Server 9

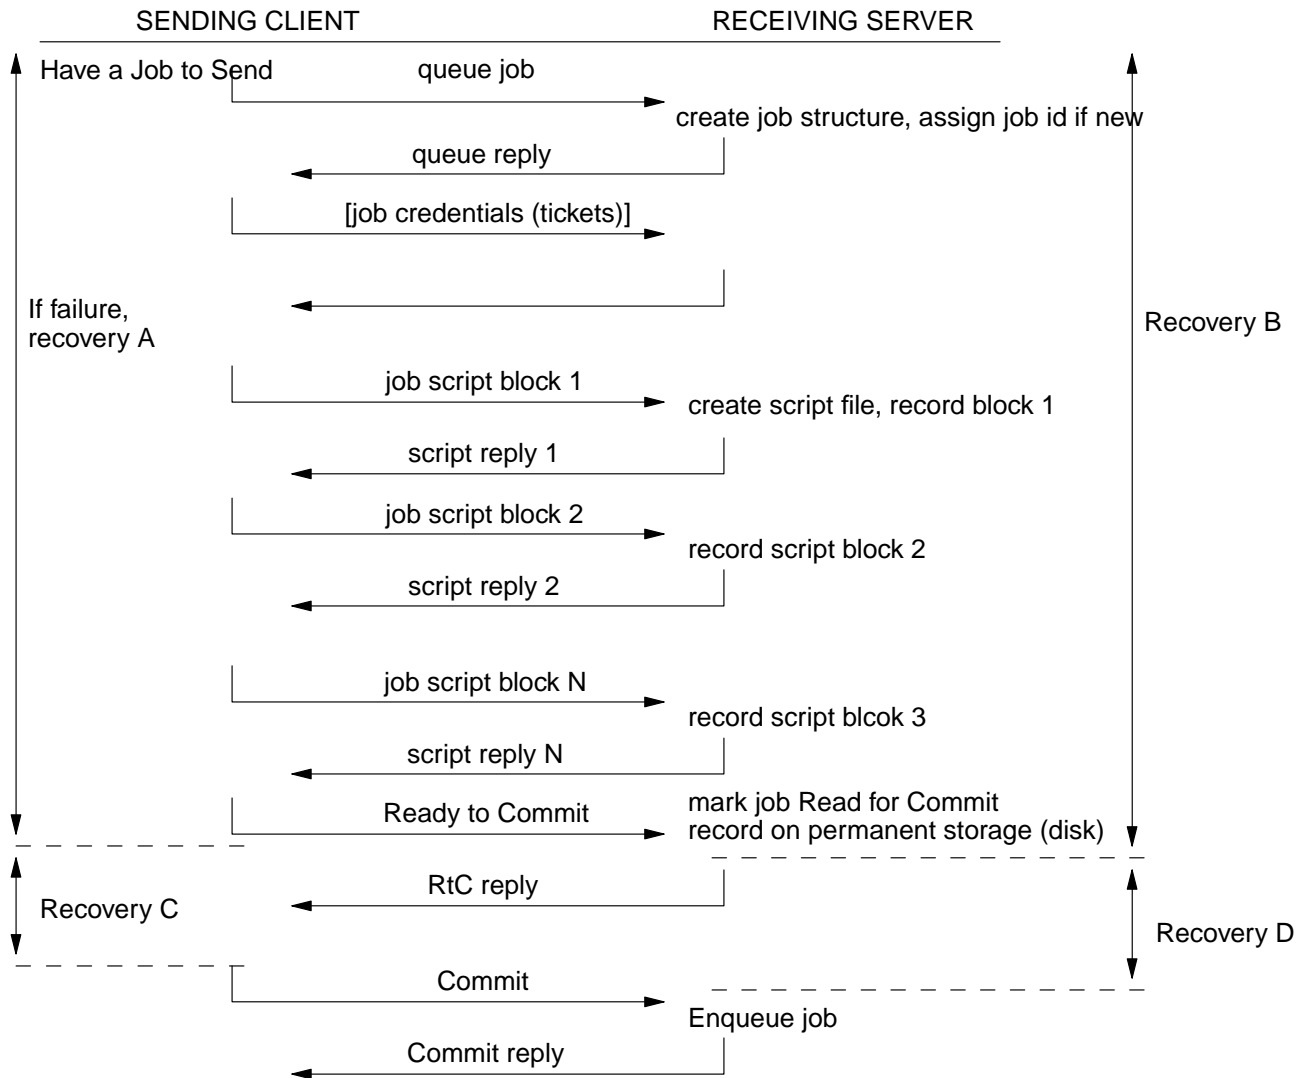
Should the receiving server crash, it should perform the following actions upon being restarted. 9

1. If the server detects a job for which a *Ready to Commit* had not been received, the server purges the job. 9



2. If the server detects a job for which the *Ready to Commit* has been received, and the sending client was a permanent process (another server), the server sends an acknowledgment of the Ready to Commit to the client. The acknowledge will include the job identifier. This will trigger the client to resend the commit sequence. In transaction terminology, this is called a *restart* message.

The send — receive — recovery steps are show in figure 11-1.



**Recovery Procedures**

- A. If sender is a server, restart from top, else abort.
- B. Discard job.
- C. Resend “Ready to Commit”.
- D. Commit (enqueue) job.

**Figure 11-2: Two Phase Commit and Error Recovery**

11.7.2. Delete Job Request

The **Delete Job** request is used by a client, such as the **qdel** command, to request that a job be deleted. The request is a *manage* request with the object type set to “job”, the name set to

|                                                                                                                                                                                                                                                                                                                                                                                                                                          |   |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| the job id, and the operation is “delete”.                                                                                                                                                                                                                                                                                                                                                                                               | 9 |
| <b>11.7.3. Hold Job Request</b>                                                                                                                                                                                                                                                                                                                                                                                                          | 9 |
| The <b>Hold Job</b> request is used by a client, such as the <b>qhold</b> command, to request that a hold be placed on a job.                                                                                                                                                                                                                                                                                                            | 9 |
| The request is a <i>manage</i> form with the object type set to “job”, the name set to the job identifier, the operation set to “set”, and the attribute list contains the Hold_Types attribute with the value being the holds to apply to the job.                                                                                                                                                                                      | 9 |
| PBS will use the request extension to pass the re-execution preference information, see <code>pbs_holdjob(3B)</code> .                                                                                                                                                                                                                                                                                                                   | 9 |
| <b>11.7.4. Modify Job Request</b>                                                                                                                                                                                                                                                                                                                                                                                                        | 9 |
| A <b>Modify Job</b> request is sent by a client such as the <b>qalter</b> command to a server to request that certain attributes of the job be altered or modified.                                                                                                                                                                                                                                                                      | 9 |
| The request is a <i>manage</i> form with the object type set to “job”, the name set to the job identifier, the operation set to “set”, and the attribute list contains each of the attributes to be modified.                                                                                                                                                                                                                            | 9 |
| <b>11.7.5. Move Job Request</b>                                                                                                                                                                                                                                                                                                                                                                                                          | 9 |
| The <b>Move Job</b> request is used by a client, such as the <b>qmove</b> command, to request that a server route the specified job to another destination.                                                                                                                                                                                                                                                                              | 9 |
| The <i>jobid</i> field specifies the job to be routed. The <i>destin</i> field specifies the destination to which the job is to be routed.                                                                                                                                                                                                                                                                                               | 9 |
| <b>11.7.6. Message Job Request</b>                                                                                                                                                                                                                                                                                                                                                                                                       | 9 |
| The <b>Message Job</b> request is used by a client, such as the <b>qmsg</b> command, to ask a server to write a message into an output file of a running job.                                                                                                                                                                                                                                                                            | 9 |
| The <i>jobid</i> field specifies the id of the running job for which the message is destined. The destination file for the message is defined by the <i>fileopt</i> ; if either the {oflg} bit and / or the {eflg} bit is set, the message will be written in the standard output and / or standard error file of the job. If neither bit is set, the message will be written to an file determined by the implementation of the server. | 9 |
| The message string will be in the <i>message</i> field.                                                                                                                                                                                                                                                                                                                                                                                  | 9 |
| <b>11.7.7. Rerun Job Request</b>                                                                                                                                                                                                                                                                                                                                                                                                         | 9 |
| The <b>Rerun Job</b> request is used by a client such as <b>qrerun</b> , to request that a server terminate and re-queue a currently running job.                                                                                                                                                                                                                                                                                        | 9 |
| The request specifies the job identifier of the running job which is to be rerun.                                                                                                                                                                                                                                                                                                                                                        | 9 |
| <b>11.7.8. Select Jobs Request</b>                                                                                                                                                                                                                                                                                                                                                                                                       | 9 |
| The <b>Select Jobs</b> request, issued by the command <code>qselect</code> , requests that the server return job identifiers for all jobs which meet the specified criteria.                                                                                                                                                                                                                                                             | 9 |
| The <i>attribute list</i> specifies the attributes which are to be used in filtering jobs. Resources, the job state (even though not a true attribute) and destination may be included in the at-                                                                                                                                                                                                                                        | 9 |

|                                                                                                                                                                                                                                                                                        |   |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| tribute set.                                                                                                                                                                                                                                                                           | 9 |
| The server replies to the client with the special reply <b>select-reply</b> . The count of jobs selected in response to the request is returned in <i>reply-auxcode</i> . The list of job identifiers selected is returned in <i>SelectReply</i> .                                     | 9 |
| <b>11.7.9. Signal Job Request</b>                                                                                                                                                                                                                                                      | 9 |
| The <b>Signal Job</b> request, used by a client such as <b>qsig</b> , requests that the server send a signal to a running batch job. The <i>jobid</i> field indicates which job is to be sent the signal. The <i>signal</i> field contains the name or value of the signal to be sent. | 9 |
| <b>11.7.10. Status Job Request</b>                                                                                                                                                                                                                                                     | 9 |
| A <b>Status Job</b> request directs a server to return status information about either a job whose job identifier is in <i>id</i> or the set of jobs in a destination given by <i>id</i> .                                                                                             | 9 |
| If the server owns the job or destination specified and the client is entitled to see the status information, the server replies with a <b>stat</b> reply. The stat reply is a list of zero or more <b>objstat</b> structures, one per job. Each objstat structure includes:           | 9 |
| - The type of object, i.e “job”.                                                                                                                                                                                                                                                       | 9 |
| - The job name                                                                                                                                                                                                                                                                         | 9 |
| - All public attributes of the job, see next paragraph.                                                                                                                                                                                                                                | 9 |
| - Optionally, a general text string for messages, implementation defined.                                                                                                                                                                                                              | 9 |
| If the status request includes a set of specified attributes, only those attributes are included in the reply.                                                                                                                                                                         | 9 |
| <b>11.7.11. Status Queue Request</b>                                                                                                                                                                                                                                                   | 9 |
| A <b>Status Queue</b> request directs the server to return status information about the queue identified in the destination specified in the request.                                                                                                                                  | 9 |
| If a queue with the specified name exists at the server, the server returns a <b>stat</b> reply containing the attribute information about that queue.                                                                                                                                 | 9 |
| If a queue name is not specified, is null, then the server will return status about all queues.                                                                                                                                                                                        | 9 |
| <b>11.7.12. Status Server Request</b>                                                                                                                                                                                                                                                  | 9 |
| A <b>Status Server</b> request directs the server to return status information about itself. The server responds with a <b>stat</b> reply.                                                                                                                                             | 9 |
| <b>11.7.13. Server Shutdown Request</b>                                                                                                                                                                                                                                                | 9 |
| The <b>Server Shutdown</b> request directs a batch server to terminate batch service and exit. The integer value is used to indicate the manner of shutdown: {immediate} or {delay}.                                                                                                   | 9 |
| <b>11.7.14. Locate Job Request</b>                                                                                                                                                                                                                                                     | 9 |
| The <b>Locate Job</b> request is used by a client to search for the current location of a job. The job identifier is specified in the request.                                                                                                                                         | 9 |
| A server who receives a Locate Job Request for a job which the server owns or knows the name of the server which owns the job, replies with a <b>locate reply</b> . If the server does not know the location, it replies with a failure code.                                          | 9 |

The locate reply is an expanded form of reply which contains the *locate-reply* field. This field contains the name of the batch server currently owning the job. 9

11.7.15. Track Job Request 9

Job tracking is an optional implementation feature in POSIX 1003.2d. Its purpose is to provide a method for recording the current location and status of jobs in a central location to facilitate job status reporting. 9

Job tracking is supported by **PBS**. When the server accepts a job for queuing, it sends a **Track Job** request to the server which originated the job and optionally to one or more "location servers" as configured by the batch system administrator. 9

When a server receives the Track Job request, that server records the current location, specified in the *location* field, of the job identified in the *jobid* field. The *hopcount* is a count of the current number of routes or hops made by the job. This is used to insure proper sequencing of the information. 9

The *jobstate* field is provided to allow for an abbreviated status following job completion. 9

11.7.16. Run Job 9

The **Run Job** request is used to direct a server to schedule a batch job immediately for execution regardless of its state, required resources, or the servers scheduling parameters. 9

11.7.17. Manage Request 9

The **Manage** request is used by the **qmgr** command (and others) to create, delete, or alter managed objects. 9

The *command* field specifies the operation: create, delete, set, or unset. The *objtype* field identifies the type of object and the *objname* field names the specific object. Attributes and associated values of the object are listed in the field *attrib*. 9

11.7.18. Register Dependent Job 9

This request is used to communicate job dependency requirements between servers. 9

The *parentid* field specifies the parent or master job identifier. The *childid* field specifies the id of the child job the server is registering with the parent. 9

The *dependtype* specifies the type of dependency. 9

The *op* field is mostly meaningful on synchronization request. It indicates if the child job has reserved its resources and is ready to run. 9

[Page intentionally left blank.]

9

**12. Tools**

9

This section describes the PBS tools. These are programs or scripts designed to make life easier for administrators and users of PBS. The existence of the PBS interface libraries provides the basis for a rich functionality. The tools give a higher level access to this functionality.

Right now only two tools exist. Both are based on Tcl so if your system doesn't have Tcl, it doesn't have these tools either. As PBS becomes more accepted and widely used, it is hoped that programs that various sites come up with will find there way into this section of the distribution (hint hint).

**12.1. Pbs\_tclsh**

9

This is a Tcl shell program which includes the same extra commands as the Tcl scheduler. These can be found in the man page for **pbs\_sched\_tcl**.

**12.2. Pbs\_wish**

9

This is a Tk window shell program with the same additional commands as *pbs\_tclsh*. This program should make it possible to write an X11 interface program to PBS.

[Page intentionally left blank.]

9

**13. Task Manger**

9

This section describes the PBS task manager functions of the Machine Oriented Miniserver (MOM). This is a new feature of PBS and is an attempt to follow the emerging "PSCHED" API. For more information on PSCHED, please see the web page: 9

<http://parallel.nas.nasa.gov/Psched/index.html> 9

A paper (which can be found on the web site above) shows the current state of the effort: "PSCHED: An API for Parallel Job/Resource Management". 9

In the past, a running job under the control of a MOM was a single POSIX session. Now, multiple sessions may be spawned, each one called a "task", and tracked separately within the job by MOM. When the job is run, one task is started with the user selected shell running the job script as the program. This is the equivalent of the single session of the past. New variables are included in the job's environment which give it information needed to call back to MOM for task management service. 9

MOM will use another tcp port to listen for task management requests from its children. This communication will be handled by a library which sends and receives messages over a tcp socket to MOM using the DIS protocol. 9



[Page intentionally left blank.]

