

PBS Professional® 11.2

Installation and Upgrade Guide



PBS Works is a division of  Altair

Altair PBS Professional 11.2 Installation & Upgrade Guide, updated 12/16/11, edited by Anne Urban

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

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

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

Copyright notice does not imply publication.

For documentation and the PBS Works forums, go to:

Web: www.pbsworks.com

For more information, contact Altair at:

Email: pbssales@altair.com

Technical Support

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

This document is proprietary information of Altair Engineering, Inc.

Table of Contents

About PBS Documentation	v
1 Concepts and Components	1
1.1 PBS Components	2
2 Pre-Installation Planning	5
2.1 Prerequisites for Running PBS	5
2.2 Single Execution System	10
2.3 Multiple Execution Systems	11
2.4 User Requirements on UNIX/Linux	12
2.5 Recommended PBS Configurations for Windows	13
2.6 Windows User Authorization	20
2.7 Windows Caveats	24
3 Installation	25
3.1 Requirements for Installing PBS Professional	25
3.2 Licensing for PBS Professional	26
3.3 Overview of Installing PBS	27
3.4 Installation on UNIX/Linux Systems	34
3.5 Installation on Windows Systems	65
3.6 Enabling Metascheduling with HPC Basic Profile	76
3.7 After Installation	78

Table of Contents

4	Licensing	81
4.1	Introduction	81
4.2	Definitions	83
4.3	Licensing Requirements	84
4.4	Configuring PBS for Licensing	84
4.5	Redundant License Servers	90
4.6	Displaying Licensing Information	91
4.7	PBS Jobs and Licensing	93
4.8	Upgrading	97
4.9	Troubleshooting Licensing Errors	98
5	Upgrading PBS Professional	101
5.1	Types of Upgrades	101
5.2	Differences from Previous Versions	102
5.3	Licensing	103
5.4	Caveats and Advice	104
5.5	Before Upgrading from 10.x or Older to 11.x or Later	106
5.6	After Upgrading	107
5.7	Introduction to Upgrading Under UNIX/Linux	107
5.8	Overlay Upgrade Under UNIX/Linux	109
5.9	Platform-Specific UNIX/Linux Overlay Upgrades	113
5.10	Migration Upgrade Under UNIX and Linux	119
5.11	Upgrade Installation Instructions	134
5.12	Upgrading Under Windows	145
6	Starting and Stopping PBS	167
6.1	Starting and Stopping PBS: UNIX and Linux	167
6.2	Starting and Stopping PBS: Windows	177
6.3	When to Restart PBS	182
6.4	Caveats and Recommendations	182
	Appendix A: License Agreement	185
	Index	195

About PBS Documentation

Where to Keep the Documentation

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

What is PBS Professional?

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

The PBS Professional Documentation

The documentation for PBS Professional includes the following:

PBS Professional Administrator's Guide:

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

PBS Professional Quick Start Guide:

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

PBS Professional Installation & Upgrade Guide:

Contains information on installing and upgrading PBS Professional.

PBS Professional User's Guide:

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

PBS Professional Programmer's Guide:

Discusses the PBS application programming interface (API).

PBS Professional Reference Guide:

Contains PBS reference material.

PBS Manual Pages:

Describe PBS commands, resources, attributes, APIs

Ordering Software and Publications

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

Document Conventions

PBS documentation uses the following typographic conventions:

abbreviation

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

command

Commands such as `qmgr` and `scp`

input

Command-line instructions

manpage (x)

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

formats

Formats

Attributes

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

Values

Keywords, instances, states, values, labels

Definitions

Terms being defined

Output

Output or example code

File contents



Chapter 1

Concepts and Components

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

Workload management systems have three primary roles:

Queuing

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

Scheduling

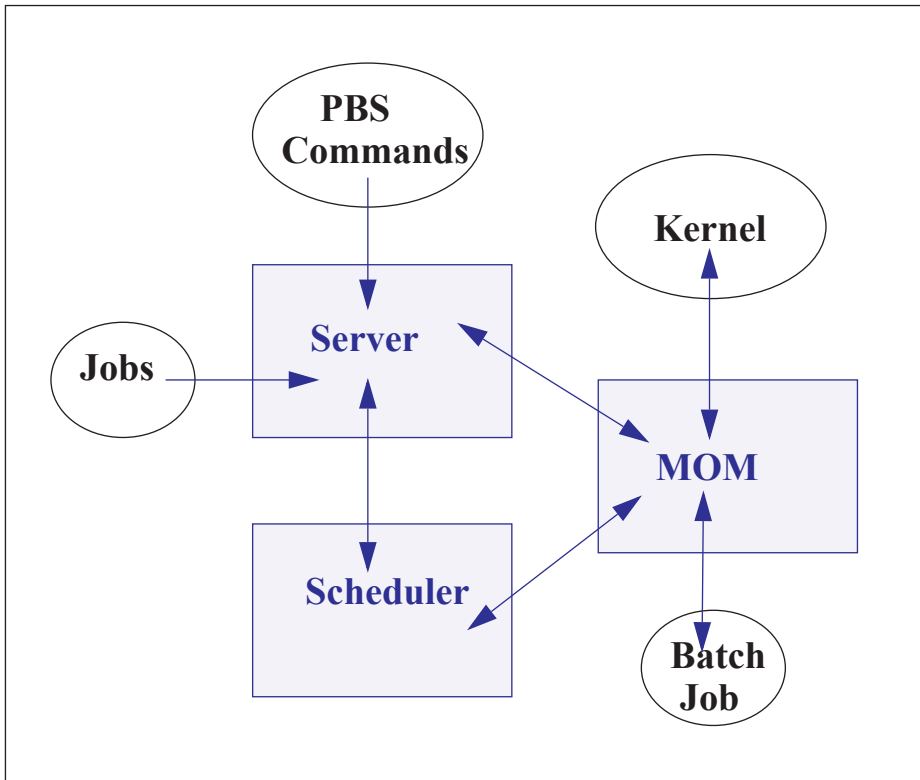
The process of selecting which jobs to run when and where, according to a predetermined policy. Sites balance competing needs and goals on the system(s) to maximize efficient use of resources (both computer time and people time).

Monitoring

The act of tracking and reserving system resources and enforcing usage policy. This covers both user-level and system-level monitoring as well as monitoring running jobs. Tools are provided to aid human monitoring of the PBS system as well.

1.1 PBS Components

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



Server

The *Server* process is the central focus for PBS. Within this document, it is generally referred to as *the Server* or by the execution name *pbs_server*. All commands and communication with the Server are via an *Internet Protocol* (IP) network. The Server's main function is to provide the basic batch services such as receiving/creating a batch job, modifying the job, protecting

the job against system crashes, and running the job. Typically there is one Server managing a given set of resources.

The server contains a licensing client which communicates with the licensing server for licensing PBS jobs.

Job Executor (MOM)

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

Scheduler

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

Commands

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

There are three classifications of commands: user commands (which any authorized user can use), operator commands, and manager (or administrator) commands. Operator and Manager commands require specific access privileges.

Chapter 2

Pre-Installation Planning

This chapter provides the information you need to know before installing PBS. Please make sure your setup meets the requirements described here, and that you take the required steps to prepare for installing PBS.

PBS is able to support a wide range of configurations. It may be installed and used to control jobs on a single system or to load balance jobs on a number of systems. It may be used to allocate vnodes of a cluster or parallel system to both parallel and serial jobs. It can also deal with a mix of these situations.

2.1 Prerequisites for Running PBS

2.1.1 Run Same Version Within Complex

Do not mix different versions of PBS within a PBS complex or across complexes. This includes both major and minor versions. All daemons and commands must be the same version of PBS Professional.

2.1.2 Resources Required by PBS

The amount of memory required by the PBS server and scheduler depends on the number of hosts and the number of jobs to be queued or running. Each host will need less than 512 bytes. The number of jobs is the important factor, since each job needs about 5 KB. The number of processors in the complex is not a factor.

2.1.2.1 Memory Required By Server Running Hooks

A PBS server executing hook scripts can consume a larger amount of memory than one not executing hook scripts. For example, a system consisting of a server and a MOM on a Linux machine or on a Windows XP machine handling 10,000 short-running jobs being submitted, modified, and moved causing execution of qsub, qalter, and movejob hooks will use around 40 MB of memory in a span of 24 hours.

2.1.2.2 Memory Required for Job History

Enabling job history requires additional memory for the server. When the server is keeping job history, it needs 8k-12k of memory per job, instead of the 5k it needs without job history. Make sure you have enough memory: multiply the number of jobs being tracked by this much memory. For example, if you are starting 100 jobs per day, and tracking history for two weeks, you're tracking 1400 jobs at a time. On average, this will require 14.3M of memory.

If the server is shut down abruptly, there is no loss of job information. However, the server will require longer to start up when keeping job history, because it must read in more information.

2.1.2.3 Amount of Memory in Complex

If the sum of all memory on all vnodes in a PBS complex is greater than 2 terabytes, then the Server (pbs_server) and Scheduler (pbs_sched) must be run on a 64-bit architecture host, using a 64-bit binary.

2.1.2.4 Adequate Space for Logfiles

PBS logging can fill up a filesystem. For customers running a large number of array jobs, we recommend that the filesystem where `$PBS_HOME` is located has at least 2 GB of free space for log files. It may also be necessary to rotate and archive log files frequently to ensure that adequate space remains available. (A typical PBS Professional complex will generate about 2 GB of log files for every 1,000,000 subjobs and/or jobs.)

2.1.2.5 Installation Disk Space

Make sure you have adequate disk space to install PBS. It is recommended to have at least 305 MB available, for installation alone.

2.1.3 Name Resolution and Network Configuration

The IP address of each machine in the complex should resolve to the fully qualified domain name for that machine. If you are using hostfiles, the FQDN should appear in `/etc/hosts`. The server must be able to look up the IP addresses for any execution host, any client host, and itself. It doesn't matter if the lookup is via `/etc/hosts` or a DNS server. Reverse name lookup of an IP address should return the fully qualified canonical hostname.

Each system hosting a PBS service (`pbs_mom`, `pbs_sched`, `pbs_server`) must allow the hosts' names to map to addresses that are accessible from the other PBS services' hosts. It may be necessary to edit `/etc/hosts` and perhaps also to disable `nsd` caching of host names in order to ensure that a host's name resolves to an address that is not on a loopback net.

The name used for the server in the `PBS_SERVER` variable in the `pbs.conf` file must not be longer than 64 characters. If the short name for the server resolves to the correct host, you can use this in `pbs.conf` as the value of `PBS_SERVER`. However, if the fully-qualified domain name is required in order to resolve to the correct host, then this must be the value of the `PBS_SERVER` variable.

On Windows, make sure the first nameserver resolves all the needed hostnames, including the server hostname and the domain controller host for active directory queries. Also, put explicit IP to hostname addresses in the `C:\windows\system32\drivers\etc\hosts` file. Otherwise your site will experience extreme slowdowns.

If you make these changes to a running PBS complex, you must then restart all the PBS daemons (services).

If you have multiple network interfaces on a single system, your `/etc/hosts` should be configured like this:

```
<IP address of public interface> host.domain.com host-pub
<IP address of private interface> host.domain.com host-priv
```

You must have the same canonical FQDN (`host.domain.com` is the same) for all local interfaces. The aliases (`host-pub` and `host-priv`) are just examples, and are not necessary.

Given that PBS is a distributed networked application, it is important that the network on which you will be deploying PBS is configured according to IETF standards. Specifically, forward and reverse name lookup should operate according to the standard.

2.1.3.1 Sockets

Some PBS processes cause network sockets to be opened between submission and execution hosts. For more information about these processes, see section 10.3.8, "Advice and Caveats", on page 634 of the PBS Professional Administrator's Guide. Make sure your network and firewalls are set up to handle sockets correctly.

2.1.3.2 Firewalls

PBS needs to be able to use any port for outgoing connections, but only specific ports for incoming connections. If you have firewalls running on the server or execution hosts, be sure to allow incoming connections on the appropriate ports for each host. By default, the PBS server and MOM daemons use ports 15001 through 15004 for incoming connections, and any port below 1024 for outgoing connections. See section 3.3.6, “Network Addresses and Ports”, on page 30 for a list of ports.

Firewall-based issues are often associated with server-MOM communication failures and messages such as 'premature end of message' in the log files.

2.1.3.3 Planning for File Access

In distributed environments it will be necessary to plan for how the users will access their input files, datasets, etc. Various options exist (such as NFS, rcp, scp, etc.). These need to be considered prior to installing PBS Professional, as such decisions can change which parameters are selected for tuning PBS. For details, see the `usecp` MOM configuration parameter in the `pbs_mom(8B)` man page and section 9.6, "Setting File Copy Mechanism" on page 574 in the PBS Professional Administrator's Guide. The impact of file location and delivery are discussed further in section 7.5 “Delivery of Output Files” on page 144 in the PBS Professional User's Guide.

Asynchronous writes to an NFS server can cause reliability problems. If using an NFS file system, mount the NFS file system synchronously (without caching.)

2.1.4 SGI Altix cpuset Feature Requires ProPack Library or Performance Suite

Customers who intend to run PBS Professional on SGI Altix systems using cpusets should note that there are strict requirements for SGI ProPack or Performance Suite (containing the cpuset API). ProPack 6 or 7 or Performance Suite 1 is required. The library is required on MOM vnodes where cpuset functionality is desired. To test whether the library is currently installed, execute the following command:

```
ls /usr/lib/libcpuset.so*
```

The PBS Professional MOM binary that supports cpusets is `pbs_mom.cpuset`.

2.1.5 Using Comprehensive System Accounting on SGI Altix

Comprehensive System Accounting (CSA) on SGI Altix requires that both the Linux job container facility and CSA support be either built into the kernel or available as a loadable module. It also requires SGI propack 6 or 7 or Performance Suite 1. See section 11.9, "Using Comprehensive System Accounting" on page 698 in the PBS Professional Administrator's Guide.

2.1.6 Requirement on Cray

PBS requires CLE 2.2 and higher on the Cray.

2.1.7 Port Requirement

The ports used by the PBS daemons must be available during the installation. See section 3.3.6, "Network Addresses and Ports", on page 30.

2.1.8 System Clocks in Sync

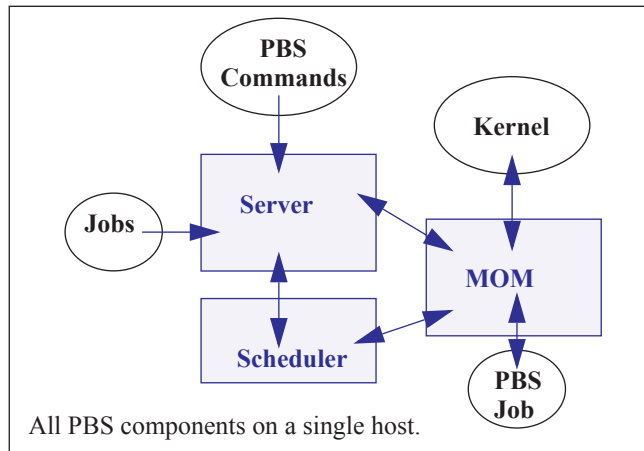
We recommend that clocks on all participating systems be in sync.

2.1.9 Windows Caveats

The PBS installer installs `vc++` redistributable binaries into the system root (`C:\Windows`) directory. See section 9.15, "Windows Caveats", on page 585 of the PBS Professional Administrator's Guide.

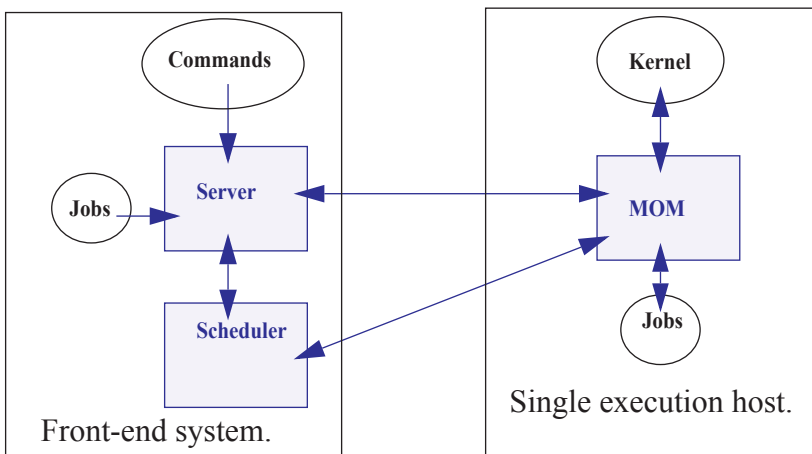
2.2 Single Execution System

If PBS is to manage a single system, all three components are installed on that same system. During installation (as discussed in the next chapter) be sure to select option 1 (all components) from the PBS Installation tool.



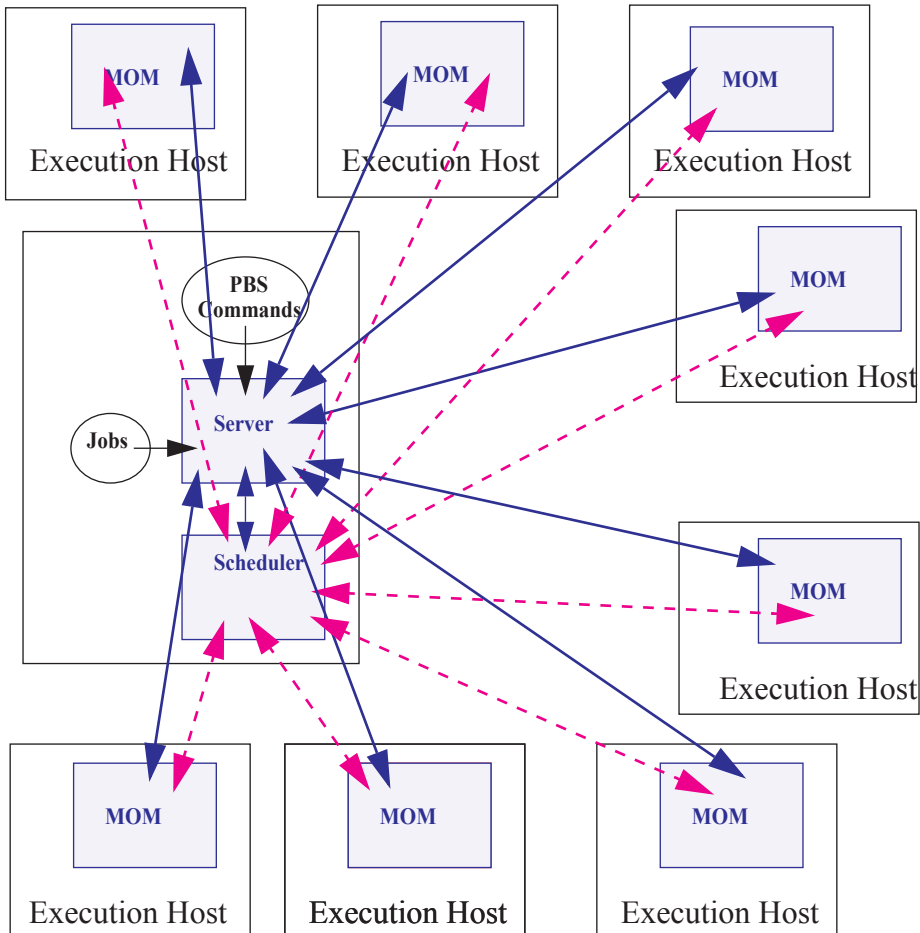
2.2.1 Single Execution System with Front-end

If you wish, the PBS Server and Scheduler (`pbs_server` and `pbs_sched`) can run on one system and jobs can execute on another.



2.3 Multiple Execution Systems

If PBS is to be installed on a collection (or complex) of systems, normally the Server (`pbs_server`) and the Scheduler (`pbs_sched`) are installed on a “front end” system (option 1 from the PBS Installation tool), and a MOM (`pbs_mom`) is installed (option 2 from the Installation tool) and run on each execution host (i.e. each system where jobs are to be executed). The following diagram illustrates this for an eight host complex.



2.4 User Requirements on UNIX/Linux

2.4.1 User Accounts

Users who will submit jobs must have accounts at the server and at each execution host.

2.4.2 UNIX User Authorization

When the user submits a job from a system other than the one on which the PBS Server is running, system-level user authorization is required. This authorization is needed for submitting the job and for PBS to return output files (see also section 7.5 “Delivery of Output Files” on page 144 in the PBS Professional User’s Guide and section 7.6 “Input/Output File Staging” on page 144 in the PBS Professional User’s Guide).

The username under which the job is to be executed is selected according to the rules listed under the “-u” option to `qsub`. The user submitting the job must be authorized to run the job under the execution user name (whether explicitly specified or not).

Such authorization is provided by any of the following methods:

1. The host on which `qsub` is run (i.e. the submission host) is trusted by the server. This permission may be granted at the system level by having the submission host as one of the entries in the server’s `host.equiv` file naming the submission host. For file delivery and file staging, the host representing the source of the file must be in the receiving host’s `host.equiv` file. Such entries require system administrator access.
2. The host on which `qsub` is run (i.e. the submission host) is explicitly trusted by the server via the user’s `.rhosts` file in his/her home directory. The `.rhosts` must contain an entry for the system from which the job is submitted, with the user name portion set to the name under which the job will run. For file delivery and file staging, the host representing the source of the file must be in the user’s `.rhosts` file on the receiving host. It is recommended to have two lines per host, one with just the “base” host name and one with the full hostname, e.g.: *host.domain.name*.
3. PBS may be configured to use the Secure Copy (`scp`) for file transfers. The administrator sets up SSH keys as described in section 9.11, "Enabling Hostbased Authentication on Linux" on page 578 in the PBS Professional Administrator’s Guide. See also section 7.5 “Delivery of Output Files” on page 144 in the PBS Professional User’s Guide.
4. User authentication may also be enabled by setting the server’s `flatuid` attribute to *True*. See the `pbs_server_attributes(7B)` man page and section 9.3.13, "Flatuid and Access", on page 568 of the PBS Professional Administrator’s Guide. Note that `flatuid`

may open a security hole in the case where a vnode has been logged into by someone impersonating a genuine user.

2.5 Recommended PBS Configurations for Windows

This section presents the recommended configuration for running PBS Professional under Windows.

2.5.1 Definitions

Active Directory

Active Directory is an implementation of LDAP directory services by Microsoft to use in Windows environments. It is a directory service used to store information about the network resources (e.g. user accounts and groups) across a domain. It was released first with Windows 2000 Server, and extended/improved under Windows Server 2003. Active Directory is fully integrated with DNS and TCP/IP; DNS is required. To be fully functional, the DNS server must support SRV resource records or service records.

Admin (Windows)

As referred to in various parts of this document, it is a user logged in from an account who is a member of any group that has full control over the local computer, domain controller, or is allowed to make domain and schema changes to the Active directory.

Administrators

A group that has built-in capabilities that give its members full control over the local system, or the domain controller host itself.

Delegation

A capability provided by Active Directory that allows granular assignment of privileges to a domain account or group. So for instance, instead of adding an account to the “Account Operators” group which might give too much access, then delegation allows giving the account read access only to all domain users and groups information. This is done via the Delegation wizard.

Domain Admin Account

It is a domain account on Windows that is a member of the “Domain Admins” group.

Domain Admins

A global group whose members are authorized to administer the domain. By default, the Domain Admins group is a member of the Administrators group on all computers that have joined a domain, including the domain controllers.

Domain User Account

It is a domain account on Windows that is a member of the “Domain Users” group.

Domain Users

A global group that, by default, includes all user accounts in a domain. When you create a user account in a domain, it is added to this group automatically.

Enterprise Admins

A group that exists only in the root domain of an Active Directory forest of domains. The group is authorized to make forest-wide changes in Active Directory, such as adding child domains.

Install Account, Installation Account

The account used by the person who installs PBS.

Schema Admins

A group that exists only in the root domain of an Active Directory forest of domains. The group is authorized to make schema changes in Active Directory.

PBS service account

The account that is used to execute the PBS daemons `pbs_server`, `pbs_mom`, `pbs_sched`, and `pbs_rshd` via the Service Control Manager on Windows. This account can have any name. The default name is *pbsadmin*.

2.5.2 Windows Configuration in a Domained Environment

2.5.2.1 Machines

- The PBS clients, server, MOMs, scheduler, and rshds must run on a set of Windows

machines networked in a single domain.

- The machines must be members of this one domain, and they must be dependent on a centralized database located on the primary/secondary domain controllers.
- The domain controllers must be running on a Server type of Windows host, using Active Directory configured in "native" mode.
- The choice of DNS must be compatible with Active Directory.
- The PBS server and scheduler must be run on a Server type of Windows machine that is not the domain controller and is running Active Directory.
- PBS must not be installed or run on a Windows machine that is serving as the domain controller (running Active Directory) to the PBS hosts.

2.5.2.2 Installation Account

- The installation account is the account from which PBS is installed. The installation account must be the only account that will be used for all steps of PBS installation including modifying configuration files, setting up failover, and so on. If any of the PBS configuration files are modified by an account that is not the installation account, permissions/ownerships of the files could be reset, rendering them inaccessible to PBS. For domained environments, the installation account must be a local account that is a member of the local Administrators group on the local computer.

2.5.2.3 The PBS Service Account

- The PBS service account is the account under which the PBS services (`pbs_server`, `pbs_mom`, `pbs_sched`, `pbs_rshd`) will run.
- This account can have any name.
- The name of the account defaults to *pbsadmin*.
- This account must exist while any PBS services are running.
- The password for this account should not be changed while PBS is running.
- Create the PBS service account before installing PBS.
- For domained environments, the PBS service account must:
 - a. be a domain account
 - b. be a member of the "Domain Users" group, and **only** this group
 - c. have "domain read" privilege to all users and groups.
- For a domained environment, delegate "read access to all users and groups information"

-
- to the PBS service account. See section 2.5.2.4, “Delegating Read Access”, on page 16.
- If the PBS service account is set up with no explicit domain read privilege, MOM may hang. This happens when XP users submit jobs from a network mapped drive without the `-o/-e` option for redirecting files. When this happens, bring up Task manager, look for a "cmd" process by the user who owned the job, and kill it. After the first cmd process is killed, you may have to look for a second one (the first one copies the output file, the second one does the error file). This should un-hang the MOM.

2.5.2.4 Delegating Read Access

- To delegate “read access to users and groups information” to the PBS service account:
 - a. On the domain controller host, bring up Active Directory Users and Computers.
 - b. Select <domain name>, right mouse click, and choose "Delegate Control". This will bring up the "Delegation of Control Wizard".
 - c. When it asks for a user or group to which to delegate control, select the name of the PBS service account.
 - d. When it asks for a task to delegate, specify "Create a custom task to delegate".
 - e. For active directory object type, select the "this folder, existing objects in this folder, and creation of objects in this folder" button.
 - f. For permissions, select “Read” and “Read All Properties”.
 - g. Exit out of Active Directory.

2.5.2.5 User Accounts

- Each user must explicitly be assigned a HomeDirectory sitting on some network path. PBS does not support a HomeDirectory that is not network-mounted. PBS currently supports network-mounted directories that are using the Windows network share facility.
- If a user was not assigned a HomeDirectory, then PBS uses `PROFILE_PATH\My Documents\PBS Pro`, where `PROFILE_PATH` could be, for example, “\Documents and Settings\username”.

2.5.2.6 User Jobs

- All users must submit and run PBS jobs using only their domain accounts (no local accounts), and domain groups. If a user has both a domain account and local account, then PBS will ensure that the job runs under the domain account.
- Each user must always supply an initial password in order to submit jobs. This is done by running the `pbs_password` command at least once to supply the password that PBS

will use to run the user's jobs.

- Access by jobs to network resources, such as a network drive, requires a password.
- All job scripts, as well as input, output, error, and intermediate files of a PBS job must reside in an NTFS directory.

2.5.2.7 Adding to Local Administrators Group

- Add the PBS service account to the local Administrators group:
`net localgroup Administrators <domain name>\<service account name> /add`

2.5.2.8 Installation Path

- The destination/installation path of PBS must be NTFS. All PBS configuration files must reside on an NTFS filesystem.

2.5.2.9 Installation

- The same PBS service account must be used in all future invocations of the install program when setting up a complex of PBS hosts.
- The install program requires the installer to supply the password for the PBS service account. This same password must be supplied to future invocations of the install program on other servers/hosts.
- The install program will enable the following rights to the PBS service account: "Create Token Object", "Replace Process Level Token", "Log On As a Service", and "Act As Part of the Operating System".
- The install program will enable Full Control permission to local "Administrators" group on the install host for all PBS-related files.
- The install program will give you a specific error if it fails to make the PBS service account be a member of the local Administrators group on the local computer. It will quit at this point, and you must go back and make the PBS service account be a member of the local Administrators group on the local computer. See section 2.5.2.7, "Adding to Local Administrators Group", on page 17. Then re-run the install program.

2.5.3 Windows Configuration in a Standalone Environment

2.5.3.1 Machines

- PBS must be run on a set of machines that are not members of any domain (workgroup

setup); no domain controllers/ Active Directory are involved.

2.5.3.2 Installation Account

- The installation account is the account from which PBS is installed. The installation account must be the only account that will be used in all aspects of PBS installation including modifying configuration files, setting up failover, and so on. If any of the PBS configuration files are modified by an account that is not the installation account, permissions/ownerships of the files could be reset, rendering them inaccessible to PBS.
- For standalone environments, the installation account must be a local account that is a member of the local Administrators group on the local computer.

2.5.3.3 The PBS service account for Standalone Environments

- The PBS service account is the account under which PBS services (`pbs_server`, `pbs_mom`, `pbs_sched`, `pbs_rshd`) will run.
- This account can have any name.
- The name of this account defaults to *pbsadmin*.
This account must exist while any PBS services are running.
- The password for this account should not be changed while PBS is running.
- Create the PBS service account before installing PBS.
- For standalone environments, the PBS service account must be a local account that is a member of the local Administrators group on the local computer. See section 2.5.3.5, “Adding to Local Administrators Group”, on page 19.
- User Accounts
- Each user should be assigned a local NTFS HomeDirectory.
- If a user was not assigned a HomeDirectory, then PBS uses `PROFILE_PATH\My Documents\PBS Pro`, where `PROFILE_PATH` could be, for example, “\Documents and Settings*username*”.
- A local account/group having the same name must exist for the user on all the execution hosts.

2.5.3.4 User Jobs

- Users must submit and run PBS jobs using their local accounts and local groups.
- Users should supply a password when submitting jobs. That password must be the same on all the execution hosts. See also section 9.13, "User Passwords", on page 582 of the

PBS Professional Administrator's Guide.

- Job intermediate, input, output, and error files must reside on an NTFS filesystem.
- For a password-ed job, users can access (within the job script) folders in a network share.
- Each user must always supply an initial password to submit jobs. This is done by running the `pbs_password` command at least once to supply the password that PBS will use to run the user's jobs.
- Access by jobs to network resources (such as a network drive), requires a password.

2.5.3.5 Adding to Local Administrators Group

- Add the PBS service account to the local Administrators group:
`net localgroup Administrators <name of PBS service account> /add`

2.5.3.6 Installation Path

- The destination/installation path of PBS must be NTFS. All PBS configuration files must reside on an NTFS filesystem.

2.5.3.7 Installation

- The same PBS service account must be used in all future invocations of the install program when setting up a complex of PBS hosts.
- The install program will require the installer to supply the password for the PBS service account. The same password must be supplied to future invocations of the install program on other Servers/hosts.
- The install program will enable the following rights to the PBS service account: "Create Token Object", "Replace Process Level Token", "Log On As a Service", and "Act As Part of the Operating System".
- The install program will enable Full Control permission to local "Administrators" group on the install host for all PBS-related files.

2.5.4 Caveats

- If you change the name of the PBS service account, you must restart the daemons on that host.

2.5.5 PBS Service Account

The PBS service account must be the same on the primary and secondary server hosts.

The PBS service account can be different on each execution host.

2.5.6 Unsupported Windows Configurations

The following Windows configurations are currently unsupported:

- PBS running on a set of Windows XP hosts that are involved in several “domains” via any trust mechanism.
- Using NIS/NIS+ for authentication on non-domain accounts.
- Using RSA SecurID module with Windows logons as a means of authenticating non-domain accounts.

2.5.7 Sample Windows Deployment Scenario

For planning and illustrative purposes, this section describes deploying PBS Professional on a complex of 20 machines networked in a single domain, with host 1 as the server host, and hosts 2 through 20 as the execution hosts.

For this configuration, the installation program is run 20 times, invoked once per host. “All” mode (server, scheduler, MOM, and `rshd`) installation is selected only on host 1, and “Execution” mode (MOM and `rshd`) installs are selected on the other 19 hosts. The PBS service account must exist in advance. This account is used for running the PBS services.

The user who runs the `install` program will supply the password for the PBS service account. The installation program will then propagate the password to the local Services Control Manager database.

A reboot of each machine is necessary at the end of each install.

2.6 Windows User Authorization

When the user submits a job from a system other than the one on which the PBS Server is running, system-level user authorization is required. This authorization is needed for submitting the job and for PBS to return output files. See section 7.5 “Delivery of Output Files” on page 144 in the PBS Professional User’s Guide and section 7.6 “Input/Output File Staging” on page 144 in the PBS Professional User’s Guide.

If running in a domained environment, then a password is also required for user authorization. See the discussion of single-signon in section 9.13, "User Passwords", on page 582 of the PBS Professional Administrator's Guide.

The username under which the job is to be executed is selected according to the rules listed under the "-u" option to `qsub`. See "qsub" on page 206 of the PBS Professional Reference Guide. The user submitting the job must be authorized to run the job under the execution user name (whether explicitly specified or not).

Such authorization is provided by either of the following methods:

1. The host on which `qsub` is run (i.e. the submission host) is trusted by the execution host. This permission may be granted at the system level by having the submission host as one of the entries in the execution host's `host.equiv` file naming the submission host. For file delivery and file staging, the host representing the source of the file must be in the receiving host's `host.equiv` file. Such entries require system administrator access.
2. The host on which `qsub` is run (i.e. the submission host) is explicitly trusted by each execution host via the user's `.rhosts` file in his/her home directory. The `.rhosts` must contain an entry for the system on which the job will execute, with the user name portion set to the name under which the job will run. For file delivery and file staging, the host representing the source of the file must be in the user's `.rhosts` file on the receiving host. It is recommended to have two lines per host, one with just the "base" host name and one with the full hostname, e.g.: *host.domain.name*.

2.6.1 Windows `hosts.equiv` File

The Windows `hosts.equiv` file determines the list of non-Administrator accounts that are allowed access to the local host, that is, the host containing this file. This file also determines whether a remote user is allowed to submit jobs to the local PBS Server, with the user on the local host being a non-Administrator account.

This file is usually: `%WINDIR%\system32\drivers\etc\hosts.equiv`.

The format of the `hosts.equiv` file is as follows:

```
[+|-] hostname username
```

'+' means enable access, whereas '-' means to disable access. If '+' or '-' is not specified, then this implies enabling of access. If only *hostname* is given, then users logged into that host are allowed access to like-named accounts on the local host. If only *username* is given, then that user has access to all accounts (except Administrator-type users) on the local host. Finally, if both *hostname* and *username* are given, then user at that host has access to like-named account on local host.

2.6.1.1 Ownership Requirements

The `hosts.equiv` file must be owned by an admin-type user or group, with write access granted to an admin-type user or group.

2.6.2 Windows User's **HOMEDIR**

Each Windows user is assumed to have a home directory (**HOMEDIR**) where his/her PBS job will initially be started. For jobs that do not have their staging and execution directories created by PBS, the home directory is also the starting location of file transfers when users specify relative path arguments to `qsub/qalter -w stagein/stageout` options.

2.6.2.1 Configuring User's **HOMEDIR**

The home directory can be configured by an Administrator by setting the user's HomeDirectory field in the user database, via the User Management Tool. It is important to include the drive letter when specifying the home directory path. The directory specified for the home folder must be accessible to the user. If the directory has incorrect permissions, PBS will be unable to run jobs for the user.

2.6.2.2 Directory Must Exist Already

You must specify an already existing directory for home folder. If you don't, the system will create it for you, but set the permissions to that which will make it inaccessible to the user.

2.6.2.3 Default Directory

If a user has not been explicitly assigned a home directory, then PBS will use this Windows-assigned default, local home directory as base location for its default home directory. More specifically, the actual home path will be:

```
[PROFILE_PATH]\My Documents\PBS Pro
```

For instance, if a *userA* has not been assigned a home directory, it will default to a local home directory of:

```
\Documents and Settings\userA\My Documents\PBS Pro
```

UserA's job will use the above path as working directory, and for jobs that do not have their staging and execution directories created by PBS, any relative pathnames in `stagein`, `stageout`, `output`, `error` file delivery will resolve to the above path.

Note that Windows can return as `PROFILE_PATH` one of the following forms:

```

\Documents and Settings\username
\Documents and Settings\username.local-hostname
\Documents and Settings\username.local-hostname.00N where N is a
  number
\Documents and Settings\username.domain-name

```

2.6.2.4 Network Mounted Home Directory

A user can be assigned a HomeDirectory that is network mounted. For instance, a user's directory can be: "\\fileserver_host\sharename". This would cause PBS to map this network path to a local drive, say G:, and allow the working directory of user's job to be on this drive. It is important that the network location (file server) for the home directory be on a Server-type of Windows machine. Workstation-type machines like Windows XP Professional have an inherent limit on the maximum number of outgoing network connections (10) which can cause PBS to fail to map or even access the user's network HomeDirectory. The net effect is the job's working directory ends up in the user's default directory:

```
PROFILE_PATH\My Documents\PBS Pro
```

If a user has been set up with a home directory network mounted, such as referencing a mapped network drive in a `HOMEDIR` path, then the user must submit jobs with a password either via `qsub -wpwd=""`, or via the single-signon feature. See “qsub” on page 206 of the PBS Professional Reference Guide and section 9.13, “User Passwords”, on page 582 of the PBS Professional Administrator’s Guide. When PBS runs the job, it will change directory to the user's home directory using his/her credential which must be unique (and passworded) when network resources are involved.

To avoid having to require passworded jobs, users must be set up to have a local home directory. Do this by accessing *Start Menu->Control Panel->Performance and Maintenance->Administrative Tools->Computer Management* (Windows XP), and selecting *System Tools->Local Users and Groups*, double clicking *Users* on the right pane, double clicking on the *username* which brings up the user properties dialog from which you can select *Profile*, and specify an input for *Local path* (Home path). Be sure to include the drive information.

2.7 Windows Caveats

2.7.1 Installation of Microsoft Redistributable Pack

The PBS installer installs the Microsoft 2005 redistributable pack. Please refer to the Microsoft documentation for further details on this package.

2.7.2 Make Sure ComSpec Environment Variable Is Set

Check that in the `pbs_environment` file, the environment variable `ComSpec` is set to `C:\WINDOWS\system32\cmd.exe`. If it is not, set it to that value:

1. Change directory:

```
cmd.admin> cd \Program Files\PBS Pro\home
```

2. Edit the `pbs_environment` file:

```
cmd.admin> edit pbs_environment
```

3. Add the following entry to the `pbs_environment` file:

```
ComSpec=C:\WINDOWS\system32\cmd.exe
```

4. Restart the MOM and the server:

```
net stop pbs_mom
net start pbs_mom
net stop pbs_server
net start pbs_server
```

2.7.2.1 Caveat

Simply setting this variable inside a job script doesn't work. The `ComSpec` variable must be set before PBS executes `cmd`. `cmd` invokes the user's submission script.

Chapter 3

Installation

As of version 11.0, PBS uses a **new** Altair license server, and you **must** use this new server.

The license server must be installed and configured before installing PBS. This chapter shows how to install PBS Professional. You should read the Release Notes and Chapter 2, "Pre-Installation Planning", on page 5 before installing the software. If you are installing on a system where PBS is already running, you probably want to preserve your configuration and your jobs during the upgrade. See Chapter 5, "Upgrading PBS Professional", on page 101.

3.1 Requirements for Installing PBS Professional

There are two steps to installing PBS. The first step is to install and configure the **new** license server package, and the second step is to install PBS itself on a selected platform. Note that if the Altair license server that is new for 11.0 is not installed and started first, PBS will **not** run properly.

Download the PBS Professional package. The PBS Professional package is available on the PBS download page at <https://secure.altair.com/UserArea/>.

The License Server package must be installed on the license server host(s) and configured before installing the PBS Professional package. PBS will not run jobs without an installed and configured license server.

The PBS server contains a licensing client. This licensing client, and therefore the PBS server, is supported on the platforms listed in the release notes.

When you run the PBS install program on platforms that don't support the licensing client, you will not be prompted for a "server type" installation option; only for "commands-only" or "execution type". The server/scheduler component in the PBS Professional package will not be included in platforms that don't support the licensing client. A list of these platforms is provided in the Release Notes.

3.1.1 Major Steps

First, the license server must be installed and configured. See section 3.2.3, "Installing the License Server", on page 27. Then you can install and configure PBS Professional. During installation of PBS Professional, you are prompted for the location of the licensing server. This sets the value of the `pbs_license_info` server attribute, so that it points to the licensing server. If this attribute is unset after PBS is installed, you must set it. See section 4.4.3.1, "Setting the Licensing Information in `pbs_license_info`", on page 87.

3.2 Licensing for PBS Professional

3.2.1 Overview of Downloading and Licensing PBS Professional

1. An administrator of the site downloads the PBS and License Server packages.
2. An administrator of the site installs the License Server package on the selected host, HostB, and configures the license server.
3. An administrator installs the PBS package on the selected host, HostA, configuring it to point to a license server on HostB. HostA may be the same as HostB.

3.2.2 Using Redundant License Servers

You may wish to set up redundant license servers. These must be set up before starting PBS. See the Altair License Management System Installation and Operations Guide, available at the Altair website, www.altair.com. After the license servers are set up, configure PBS to use them. See section 4.5, "Redundant License Servers", on page 90.

3.2.3 Installing the License Server

For instructions on downloading and installing the Altair License Manager, see the Altair License Management System Installation and Operations Guide, available at the Altair website, www.altair.com.

3.3 Overview of Installing PBS

3.3.1 Major Steps in Installing PBS

1. Create required accounts. See section 3.4.1, “Create Required Accounts”, on page 34 and section 3.5.4, “Pre-installation Configuration”, on page 66.
2. Download the PBS software
3. Install PBS on the server host
4. Install PBS on any client hosts
5. Start PBS on the server host by running the PBS start script. See section 6.1, “Starting and Stopping PBS: UNIX and Linux”, on page 167.
6. Install PBS on each execution host, choosing to start PBS there.
7. Using the `qmgr` command, define the vnodes that the server will manage. See section 3.7.1, “Creating or Modifying Vnodes”, on page 78.

3.3.2 Getting the Software

The PBS software can be installed from the PBS CD-ROM or downloaded from the User Login Area of the PBS website (<http://www.pbsworks.com>). The installation procedure is slightly different depending on the distribution source. However, the basic steps of PBS installation are:

1. Prepare distribution media
2. Download or extract the software
3. Install the software
 - a. Set the licensing information on the PBS server

3.3.3 Versioned Installation

Except on Windows, the PBS binaries are installed in a versioned directory, which is created by the installation program. The PBS information files are not installed in a versioned directory.

3.3.4 Setting PBS License Server Location

During installation of the server, you will be prompted for the location of the Altair license server. You must use an Altair server; you cannot use a FLEX server. See section 4.4.2, “Server Licensing Attributes”, on page 85. The installer will display the following:

```
PBS Professional version 9.0 and later is licensed via the Altair License
  Manager.
```

```
The Altair License Manager can be downloaded from:
```

```
http://www.pbspro.com/UserArea/Software/
```

```
For more information, please refer to the PBS Professional Administrator's
  Guide, or contact pbssupport@altair.com.
```

```
Continue with the installation? (y or n)
```

The PBS install program will prompt you for the location of the license server:

```
Please enter the list of Altair License file location(s), in a space-
  separated list of entries of the form:
```

```
<port>@<host>
```

```
@<host>
```

```
<license-file-path>
```

The install program then lists some examples:

```
Examples:
```

```
@fest
```

```
27100@aspasia
```

```
@perikles 27000@aspasia
```

```
@127.3.4.5
```

```
/usr/local/altair/security/altair_lic.dat
```

```
Enter License File Location(s): <>
```

NOTE: The “Enter License File Location(s):” is the actual prompt; the lines above it are informational.

The install program will take care of setting the `pbs_license_info` attribute to the location entered.

3.3.4.1 Caveats

If you do not input any entry for the "Enter License File Location(s)" prompt, then the install program will not initialize the `pbs_license_info` attribute. This means the `pbs_license_info` attribute value is left as is, which could be set to some previous value or unset. It is usually set to some previous value when doing an overlay or migration upgrade.

If the license server location is incorrectly initialized (e.g. the host name or port number is incorrect), PBS may not be able to pinpoint the misconfiguration as the cause of the failure to reach a license server. The PBS server's first attempt to contact the license server will result in the following message on the server's log file:

```
"unable to connect to license server at host <H>, port <P>"
```

If the license server location is set to a FLEX server, PBS may encounter a problem requiring you to reboot the server's host. See section 14.1, "Server Host Bogs Down After Startup", on page 757 of the PBS Professional Administrator's Guide.

3.3.5 Choosing Installation Type

You can choose the tasks for each machine in a PBS complex. These tasks are:

- Host on which server and scheduler run. This uses the server/scheduler package described below.
- Host which runs MOM, i.e. run users' tasks. This uses the execution package described in section 3.3.5.2, "Execution Package", on page 30.
- Host from which users submit jobs. This uses the client commands package described in section 3.3.5.3, "Client Commands", on page 30.

You can choose to run all PBS components, and submit jobs, and run jobs, on a single machine, or you can divide up the tasks.

A typical cluster of computers has a front-end system which usually manages the whole cluster. Most sites install the PBS Server and Scheduler on this front-end system, but not the MOM, as most sites tend *not* to want to run batch jobs on the front-end host. The MOM is then installed on each execution host within the cluster.

You need to decide what kind of PBS installation you want for each machine in your complex. During installation, there are three options:

1. Server/scheduler package, execution package, and client commands
2. Execution package and client commands
3. Client commands

3.3.5.1 Server/Scheduler Package

Install this package on the machine(s) which will run the PBS server. If you are going to use failover, install this package (option 1) on both server hosts. This package includes the server, the scheduler, a MOM, and the commands.

If you are going to run all the PBS components on a single host, install the Server package (option 1).

3.3.5.2 Execution Package

If you are going to have a complex of machines, you need to install the execution package (option 2) on all the execution hosts in the complex. This package includes the MOM and the client commands.

3.3.5.3 Client Commands

The client package (option 3) is for hosts which will not be used for execution but need to have access to PBS. It contains the commands, the GUIs, and the man pages. This allows users to submit jobs, and to check the status of jobs, queues, and any PBS Servers.

3.3.6 Network Addresses and Ports

PBS daemons listen for inbound connections at specific network ports. These ports have defaults, but can be configured if necessary. PBS daemons use any ports numbered less than 1024 for outbound communication. For PBS daemon-to-daemon communication over TCP, the originating daemon will request a privileged port for its end of the communication.

PBS makes use of fully qualified host names for identifying the jobs and their location. A PBS installation is known by the host name on which the Server is running. The canonical host name is used to authenticate messages, and is taken from the primary name field, `h_name`, in the structure returned by the library call `gethostbyaddr ()`. According to the IETF RFCs, this name must be fully qualified and consistent for any IP address assigned to that host.

Port numbers can be set via `/etc/services`, the command line, or in `pbs.conf`. If not set by any of these means, they will be set to the default values. The PBS components and the commands will attempt to use the system `services` file to identify the standard port numbers to use for communication. If the port number for a PBS service can't be found in the system file, a default value for that service will be used. Each daemon has startup options for setting port numbers. In the PBS Professional Reference Guide, see “`pbs_mom`” on page 56, “`pbs_sched`” on page 88, and “`pbs_server`” on page 90.

The table below lists the default port numbers for PBS services:

Table 3-1: Ports Used by PBS Services

Daemon Listening at Port	Port Number	Protocol	Type of Communication
<code>pbs_server</code>	<i>15001</i>	TCP	All TCP communication to server
<code>pbs_server</code>	<i>15001</i>	UDP	Server to MOM via RPP
<code>pbs_mom</code>	<i>15002</i>	TCP	MOM to/from Server
<code>pbs_resmon</code>	<i>15003</i>	TCP	MOM resource requests
<code>pbs_resmon</code>	<i>15003</i>	UDP	MOM resource requests
<code>pbs_sched</code>	<i>15004</i>	TCP	PBS Scheduler
<code>pbs_mom_globus</code>	<i>15005</i>	TCP	MOM Globus
<code>pbs_mom_globus</code>	<i>15006</i>	TCP	MOM Globus resource requests
<code>pbs_mom_globus</code>	<i>15006</i>	UDP	MOM Globus resource requests
<code>pbs_datastore</code>	<i>15007</i>	proprietary	PBS information storage and retrieval
License server	<i>6200</i>		

The scheduler uses any privileged port (less than 1024) as the outgoing port to talk to the Server.

Under UNIX, the `services` file is named `/etc/services`.

Under Windows, it is named `%WINDIR%\system32\drivers\etc\services`.

The port numbers listed are the default numbers used by PBS. If you change them, be careful to use the same numbers on all systems. The port number for `pbs_resmon` must be one higher than for `pbs_mom`.

Communication between the scheduler and server is via TCP.

3.3.7 Terminology and Conventions

3.3.7.1 Version Number

M.N.P.S refers to the PBS Professional version number. The version number has this form:

`<major>.<minor>.<patch>.<sequence>`

Example:

M: Major release number: 11

N: Minor release number: 0

P: Patch number: 0

S: Sequence number: 08765

In this example, the version number is 10.0.0.08765.

3.3.7.2 Pathname Conventions

The term `PBS_HOME` refers to the location where the daemon/service configuration files, accounting logs, etc. are installed.

The term `PBS_EXEC` refers to the location where the executable programs are installed. `PBS_EXEC` includes the directory named with the M.N.P.S version number.

`<PBS_EXEC_BASE>` is used to reference the base path of the directory into which `PBS_EXEC` is being installed. It excludes the “version” subdirectory. For example, if `PBS_EXEC` is being installed into `/opt/pbs/10.1.0.12345`, then `<PBS_EXEC_BASE>` refers to `/opt/pbs`.

3.3.8 Default Installation Locations

The installation program installs the various PBS components into specific locations on the system. The installation program allows you to override these default locations if you wish. The locations are written to the `pbs.conf` file created by the installation process. For details see section 13.1, “The PBS Configuration File”, on page 735 of the PBS Professional Administrator’s Guide. The PBS package will install into the default directory unless you specify otherwise.

3.3.8.1 Location of PBS_HOME

The default installation directories for PBS_HOME are determined by the operating system being used. During installation, if `pbs.conf` or the administrator specifies the location of PBS_HOME, PBS_HOME will be put there.

3.3.8.2 Location of PBS_EXEC

The default installation directories for PBS_EXEC are versioned. If the administrator overrides the default installation directory, the new directory must also be correctly versioned. If it is not, the version will be appended to the new directory. For example, if the administrator chooses `/usr/local/pbs`, the package is installed into `/usr/local/pbs/M.N.P.S`.

3.3.8.3 Default Installation Locations

Default installation directories are shown in the following table.

Table 3-2: Default Installation Locations

OS	Location of PBS_HOME	Location of PBS_EXEC
AIX	<code>/usr/local/spool/PBS</code>	<code>/opt/pbs/M.N.P.S</code>
HP-UX	<code>/usr/spool/PBS</code>	<code>/opt/pbs/M.N.P.S</code>
Linux	<code>/var/spool/PBS</code>	<code>/opt/pbs/M.N.P.S</code>
Solaris	<code>/usr/spool/PBS</code>	<code>/opt/pbs/M.N.P.S</code>
Windows (32-bit)	<code>C:\Program Files\PBS Pro\home</code>	<code>C:\Program Files\PBS Pro\exec</code>
Windows (64-bit)	<code>C:\Program Files (x86)\PBS Pro\home</code>	<code>C:\Program Files (x86)\PBS Pro\exec</code>

3.3.9 Instructions by Platform

The procedure for installing PBS is the same on most platforms. Some platforms have a few minor differences, and some require special instructions. You will find differences and instructions here:

Table 3-3: Installation Steps by Platform

Platform/Arrangement	Instructions
Most UNIX/Linux	section 3.4, “Installation on UNIX/Linux Systems”, on page 34
Cray	section 3.4.10, “Installing PBS on the Cray XT”, on page 56
SGI Altix	section 3.4.8, “Installing on an Altix Running SuSE”, on page 47
SGI ICE/XE	section 3.4.9, “Installing PBS on the SGI ICE/XE”, on page 49
AIX (Minor difference)	section 3.4.7.1, “Notes About AIX”, on page 44
Solaris (Minor difference)	section 3.4.7.2, “Notes About Solaris”, on page 46
Windows	section 3.5, “Installation on Windows Systems”, on page 65

3.4 Installation on UNIX/Linux Systems

3.4.1 Create Required Accounts

Before you install PBS, you must create the accounts that PBS requires. Create the data service user account, with the following characteristics:

- Non-root account
- Account is enabled
- If you are using failover, the UID of this account must be the same on both primary and secondary server hosts
- We recommend that the account is called *pbsdata*. The installer looks for an account

called *pbsdata*. If this account exists, the installer does not need to prompt for a user-name, and can install silently.

- Root must be able to `su` to the data service account and run commands as that user. Do not add lines such as `'exec bash'` to the `.profile` of the data service account. If you want to use `bash` or similar, set this in the `/etc/passwd` file, via the OS tools for user management.
- The data service account must have a home directory.

3.4.2 Permissions

The location for the installation of the PBS Professional software binaries (`PBS_EXEC`) and private directories (`PBS_HOME`) must be owned and writable by root, and must not be writable by other users.

3.4.3 Installation Programs

The PBS install script uses the native package installer provided with each system.

Some systems' installation programs (e.g. Solaris `pkgadd`) will ask you to confirm that it is acceptable to install `setuid/setgid` programs as well as to run installation sub-programs as root. You should answer `yes` or `y` to either of these questions, if asked.

3.4.4 If You Have Multiple Execution Hosts

Instead of running the installer by hand on each machine, you can use a command such as `pdsh` with NFS-mounted common directories on all hosts to distribute the installation to each host. The general form to distribute commands to a large number of hosts is

```
# for hosts node001-005
pdsh -w node[001-005] command
```

When using the PBS `INSTALL` command you can make a simple text file called “answers” with the answers to the installation prompts, and redirect them into the `INSTALL` script:

```
./INSTALL < answers
```

where `answers` contains the answers for a MOM-only installation:

```
<newline>
<newline>
2
y
<server name>
y
```

To feed your file of answers to the install script:

```
pdsh -w node[001-005] ./INSTALL < answers
```

3.4.5 Systems With Prior Versions of PBS

It is important to preserve your configuration information, and to handle existing jobs correctly. Please read Chapter 5, "Upgrading PBS Professional", on page 101.

If a prior version of PBS Professional already exists, the installation procedure will not remove it, except for Solaris. See section 3.4.7.2.i, "Automatic Uninstallation of Prior Packages on Solaris", on page 46. The installation program allows you to choose to switch to the new version of PBS.

Future versions installed into a versioned `PBS_EXEC` directory may co-reside with the other versions also installed into a versioned directory. The Administrator may uninstall the older package at any future time.

On Cray XT systems only, if the directory `/opt/modulefiles` exists, a module description file is installed in `/opt/modulefiles/pbs/M.N.P.S`. See section 3.4.10, "Installing PBS on the Cray XT", on page 56.

3.4.5.1 Which Version To Run

If a prior version of PBS is found while installing the new version of PBS, the installation program asks whether to switch to the new version of PBS. If no prior version of PBS Professional was found, the install script does not ask whether to switch to the new version.

3.4.5.1.i Switching To New Version Now

If you choose to switch to the new version, the installation program updates configuration files including `/etc/pbs.conf`. It also moves the pre-existing `/etc/pbs.conf` to `/etc/pbs.conf.pre.M.N.P.S`.

When PBS is switched to the newly installed version, a symbolic link named `<PBS_EXEC_BASE>/default` is created to point to the location of the new version. For example, `/opt/pbs/default -> /opt/pbs/10.0.0.4567`.

3.4.5.1.ii Switching To New Version Later

If you choose not to switch to the new version of PBS at the time of installation, you can switch later. To switch to the new version of PBS, run the new `pbs_postinstall` script in the newly-created directory for the version to which you are switching:

```
<PBS_EXEC_BASE>/<new M.N.P.S>/etc/pbs_postinstall
```

For example, if you are switching to version 10.0.0:

```
/opt/pbs/10.0.0.08470/etc/pbs_postinstall
```

Running `pbs_postinstall` updates `/etc/pbs.conf` with the new information and establishes `PBS_EXEC/default` as a symbolic link to the new `PBS_EXEC`, which is `/opt/pbs/10.0.0.08470` here.

3.4.5.2 Automatic Archiving and User Path Preservation

If your prior version of PBS was installed in the default location, the installation will automatically create links so that the old default paths will continue to work. Note that the new default paths are different, and users should update their paths anyway. However, if the prior version of PBS was not installed in the default location, the installation will not create links to the old paths. You may choose to do so by hand. Please read the following subsections.

3.4.5.2.i Archiving Prior Version

Older PBS Professional packages may be archived to prevent users from accidentally running the old version of commands. The installation program archives the old version of PBS if the following conditions are true:

- The prior `PBS_EXEC` directory was any of:
 - `/usr/pbs`
 - `/opt/pbs`
 - `/usr/local/pbs`
- The directory exists when the administrator switches to the new version of PBS

This is what happens when the installation program archives the old version of PBS:

1. A new directory called `PBS_EXEC/pre.M.N.P.S` is created, for example “`/usr/pbs/pre.10.0.0.4567`”.
2. The `bin`, `etc`, `include`, `lib`, `man`, `sbin`, `tcltk` and `unsupported` directories are moved into the newly created directory.

3.4.5.2.ii Automatic Preservation of User Paths

If the PBS installation program was able to archive the older version of PBS, it also creates symbolic links where possible so that user paths are unchanged. The following symbolic links are created:

```
OLD_PBS_EXEC/bin -> /<PBS_EXEC_BASE>/default/bin
OLD_PBS_EXEC/sbin -> /<PBS_EXEC_BASE>/default/sbin
OLD_PBS_EXEC/etc -> /<PBS_EXEC_BASE>/default/etc
```

3.4.5.2.iii Re-creating User Paths By Hand

If uninstalling the prior package removes or relocates the prior `PBS_EXEC` directory which was pointed to by users’ `PATH` variable, you may wish to recreate that directory and create under it the following symbolic links to enable users to access the new commands via their current `PATH`:

```
OLD_PBS_EXEC/bin -> <PBS_EXEC_BASE>/default/bin
OLD_PBS_EXEC/sbin -> <PBS_EXEC_BASE>/default/sbin
```

For example if the old location was `/usr/pbs_bin`, create the link `/usr/pbs_bin/bin -> /opt/pbs/default/bin`.

3.4.5.2.iv Updating User Paths

Users should set their path to include `<PBS_EXEC_BASE>/default/bin`. For example, if `<PBS_EXEC_BASE>` is `/opt/pbs`, by including `/opt/pbs/default/bin`, users will have the current version of PBS in their path.

3.4.5.3 Uninstalling Old Versions

Future versions installed into a versioned `PBS_EXEC` directory may co-reside with the other versions also installed into a versioned directory. You may uninstall the older package at any future time.

When you wish to remove a prior version which was relocated by the new install, some package managers (such as `rpm`) may be unable to remove the relocated directory. In this case, you should remove the old directory by hand. The prior version of the `/etc/pbs.conf` file was saved as `/etc/pbs.conf.pre.M.N.P.S`; remove this as well.

3.4.5.3.i Caveats

If POE, Intel MPI, HP MPI, MPICH, etc., were wrapped, you must make sure that those are unwrapped before uninstalling PBS.

On UNIX, when PBS is uninstalled, files in `PBS_HOME/datastore` may remain.

3.4.5.4 Reverting To An Older Version of PBS

To revert to a prior, still installed, version of PBS, follow these steps:

1. Shut down PBS as described in section 6.1, “Starting and Stopping PBS: UNIX and Linux”, on page 167.
2. Change the `<PBS_EXEC_BASE>/default` link to point to the desired version on the Server node and on each execution node.
3. Restart PBS as described in section 6.1, “Starting and Stopping PBS: UNIX and Linux”, on page 167.

3.4.5.4.i Dangers in Reverting

Note that reverting may remove added capabilities and may cause the loss of certain job, queue, or Server attributes that were introduced in the newer version. If the newer version introduces incompatibilities which would prevent reverting back to a prior version, this will be documented in the release notes for the newer version. Please review the release notes prior to reverting.

3.4.5.4.ii Reverting on Solaris

On Solaris system, the prior package had to be removed prior to installation of the new package; therefore the new package must be uninstalled and the prior version must be reinstalled in order to revert back to that version. See section 3.4.7.2, “Notes About Solaris”, on page 46.

3.4.6 Installation Procedure (Most Systems)

3.4.6.1 Before Installing

Please do not jump straight to this section in your reading. Before downloading and installing PBS, please make sure that you have read the following and taken any required steps:

- Prerequisites: All of Section 2.1, "Prerequisites for Running PBS", Section 2.2, "Single Execution System", Section 2.3, "Multiple Execution Systems", and Section 2.4, "User Requirements on UNIX/Linux" and their subsections.
- Pre-installation steps: All of Section 3.1, "Requirements for Installing PBS Professional", Section 3.2, "Licensing for PBS Professional", and Section 3.3, "Overview of Installing

PBS", and their subsections.

- Installation steps: Section 3.4.1, "Create Required Accounts", Section 3.4.2, "Permissions", Section 3.4.3, "Installation Programs", Section 3.4.4, "If You Have Multiple Execution Hosts", and Section 3.4.5, "Systems With Prior Versions of PBS", and their subsections.

3.4.6.2 Where to Find Information

Follow the instructions below for downloading or reading the CD-ROM. If you are installing PBS on a standalone system, follow the instructions in section 3.4.6.5, "Installing on a Standalone Machine", on page 43. If you are installing on a cluster, follow the instructions in section 3.4.6.6, "Installing on a UNIX/Linux Cluster", on page 44. An example of running the installer via ./INSTALL is shown in section 3.4.6.4, "Example of Running the Installer", on page 41.

3.4.6.3 Download/CDROM Instructions

3.4.6.3.i CD-ROM

If installing from the PBS CD-ROM:

1. Insert the PBS CD into the system CD-ROM drive
2. Mount the CD-ROM device (if needed)

```
mount /cdrom
```

3. Change directory to the distribution directory.

```
cd /cdrom/PBSPro_<M.N.P.S>
```

3.4.6.3.ii Download

1. Download the PBS tar.gz package to /tmp
2. Change directory to /tmp:

```
cd /tmp
```

3. Extract the tarfile:

```
tar zxvf PBSPro_<M.N.P.S>-linux26_i686.tar.gz
```

3.4.6.4 Example of Running the Installer

The installation process is launched by executing the `INSTALL` command, as shown below. `<M.N.P.S>` is a placeholder for the actual version number.

```
./INSTALL
Installation of PBS
The following directory will be the root of the installation. Several
subdirectories will be created if they don't already exist: bin, sbin,
lib, man and include.
Execution directory? [/opt/pbs/<M.N.P.S>]
PBS needs to have a private directory (referred to as "PBS_HOME" in the
documentation) where it can permanently store information.
Home directory? [/usr/spool/PBS]
/usr/spool/PBS does not exist, I'll make it...done
[ Description of the different configuration options ]
PBS Installation:
  1. Server, execution and commands
  2. Execution only
  3. Commands
```

Next, the installation program will proceed to extract and install the PBS package(s) that you selected above. The process should look similar to the example below.

```
## Installing part 1 of 1.
/etc/init.d/pbs
```

(listing of files not shown for brevity)

```
## Executing postinstall script.
*** PBS Installation Summary
***
*** PBS Server has been installed in /opt/pbs/<M.N.P.S>/sbin
*** PBS commands have been installed in /opt/pbs/<M.N.P.S>/bin
***
*** This host has the PBS Server installed, so
*** the PBS commands will use the local server.
*** The PBS command server host is mars
***
*** PBS MOM has been installed in /opt/pbs/<M.N.P.S>/sbin.
*** PBS Scheduler has been installed in /opt/pbs/<M.N.P.S>/sbin.
***
Installation of <pbs64> was successful.
```

3.4.6.5 Installing on a Standalone Machine

The following example shows an installation on a single host on which all PBS components will run, and from which users will also submit jobs. The process may vary depending on the native package installer on each system.

1. Change to the directory where PBS was put when extracted:

```
cd PBSPro_<M.N.P.S>
```

2. Execute INSTALL.

```
./INSTALL
```

See section 3.4.6.4, “Example of Running the Installer”, on page 41.

Choose option 1.

Set the server’s `pbs_license_info` attribute by giving the location of the license server.

See section 4.4.3.1, “Setting the Licensing Information in `pbs_license_info`”, on page 87.

3. Start PBS:

```
/etc/init.d/pbs start
```

4. Check to see that the server, scheduler and MOM daemons are running:

```
ps -ef | grep pbs
```

You should see that there are three daemons running: `pbs_mom`, `pbs_server`,
`pbs_sched`

5. Optionally re-create user paths. See section 3.4.5.2.iii, “Re-creating User Paths By Hand”, on page 38.

6. Test that a normal user can submit a job:

```
echo "sleep 60" | /opt/pbs/default/bin/qsub
```

This will submit a job in the 'workq' queue because it is the default queue defined within `qmgr`

7. Verify that user paths have been preserved. See section 3.4.5.2, “Automatic Archiving and User Path Preservation”, on page 37.

```
echo "sleep 60" | <old user path>/bin/qsub
```

8. Verify that new user path has been created:

```
echo "sleep 60" | /opt/pbs/default/bin/qsub
```

9. Verify that the jobs are running:

```
/opt/pbs/default/bin/qstat -a
```

10. Verify that you are running the correct version:

```
/opt/pbs/default/bin/qstat --version
```

3.4.6.6 Installing on a UNIX/Linux Cluster

You may or may not want to run batch jobs on the server/scheduler host. First, install and start PBS on the Server host. Then install PBS on each execution host. Follow these steps:

1. Install PBS on the cluster's front-end machine. See section 3.4.6.4, "Example of Running the Installer", on page 41.
2. When asked if you want to start PBS, enter "no".
3. If you do *not* want to run batch jobs on the front-end host:

Edit the newly installed `/etc/pbs.conf` file. Set `PBS_START_MOM=0`.

4. Start PBS on the Server machine by running the PBS startup script. The location of the script varies depending on system type. See section 6.1, "Starting and Stopping PBS: UNIX and Linux", on page 167.
5. Install PBS on each execution host. Do this by running the `INSTALL` program on each host, selecting the execution package only (option 2).

When prompted if you wish to start PBS on that host, enter **yes**.

If all execution hosts are identical, you could install one of the execution hosts, and then distribute the installation to other hosts via a program such as `rdist`, or via `tar` plus `scp/rcp`. See section 3.4.4, "If You Have Multiple Execution Hosts", on page 35.

6. Using the `qmgr` command, define the vnodes that the server will manage. See section 3.7.1, "Creating or Modifying Vnodes", on page 78.

3.4.7 Minor Differences

3.4.7.1 Notes About AIX

When you download the AIX package, you will get two versions of `pbs_mom`. One is for using with the HPS switch, and the other is the standard `pbs_mom`. The installer automatically looks for the HPS switch, and if it finds it, installs the version of `pbs_mom` that manages the switch. If the installer doesn't find the switch, it installs the standard `pbs_mom`.

3.4.7.1.i Detecting High Performance Switches

During installation, the PBS install script looks for the presence of either an InfiniBand switch or an HPS. The script prints messages to stdout.

While checking for InfiniBand and HPS:

```
*** Checking for High Performance Network
```

If the install script detects neither an InfiniBand switch nor an HPS:

```
*** INFO: No High Performance Network adaptors detected
*** It appears no High Performance Network is available
*** If this is incorrect, please see the PBS
*** Administrator's Guide to configure PBS for
*** jobs to be able use the High Performance Network
```

If the install script finds an HPS:

```
*** HPS found!
```

If the install script finds InfiniBand:

```
*** InfiniBand found!
```

The script sets the default transfer method to User Space (US) mode.

- If the script has found an HPS switch:
*** setting MP_EUILIB=us
- If the script has found an InfiniBand switch:
*** setting MP_EUILIB=us
*** setting MP_DEVTTYPE=ib

At startup the PBS MOM logs the InfiniBand adapters it has detected:

```
02/13/2009 10:34:13;0002;pbs_mom;Svr;pbs_mom;detected 4 InfiniBand
adapters connected to 2 networks using 512 windows
```

The value for windows shown in the log message above appears in the output of the `pbsnodes` command as the following:

```
resources_available.netwins=512
```

This value should not be changed, via either `qmgr` or MOM configuration files.

3.4.7.1.ii Failure to Detect InfiniBand Switch

If the install script fails to detect an InfiniBand switch, you can install the InfiniBand-configured MOM by hand:

```
cp PBS_EXEC/sbin/pbs_mom PBS_EXEC/sbin/pbs_mom.standard
cp PBS_EXEC/sbin/pbs_mom.ib PBS_EXEC/sbin/pbs_mom
```

3.4.7.2 Notes About Solaris

During installation on Solaris:

- You must specify the location of `PBS_EXEC` to the package manager
- The package manager uninstalls any existing versions of PBS during installation

3.4.7.2.i Automatic Uninstallation of Prior Packages on Solaris

If a prior version of PBS which was installed via the Solaris package manager (`pkgadd`) exists, it is removed by the install procedure. The install procedure requests permission to remove the prior version. If you do not give permission, the installation is abandoned.

The following messages are displayed:

```
There is an existing PBS installed, pbs64.
```

```
The Solaris package manager only allows one version to be installed.
```

```
Do you wish to remove the existing package and continue with the install  
(y/n)?:
```

If you answer “y”, this is the response:

```
Removing package pbs64
```

If you answer “n”, this is the response:

```
Abandoning the install.
```

3.4.7.2.ii Requirements for Installation on Solaris

Installing PBS on Solaris requires you to specify the location of `PBS_EXEC`. The Solaris package manager does not automatically install PBS into `/opt/pbs/M.N.P.S.` The package manager will ask for the location of the directory into which to install the binaries, twice, and you must make sure that the answers match perfectly.

The package manager ensures that M.N.P.S is at the end of the path that you specify. If you enter just `<PBS_EXEC_BASE>`, whether this is the default or not, it will append the M.N.P.S. If you enter an incorrect M.N.P.S, the correct one will be appended, after the incorrect one. So if the version is 10.0.0.12345, and you accidentally enter `/usr/local/100.0.0.12345`, you'll get `/usr/local/100.0.0.12345/10.0.0.12345`.

1. In the first request, the installation script offers a default directory. In this example, the version is 10.0.0.82971.

```
Execution directory? [/opt/pbs/10.0.0.82971]
```

You can do one of the following:

- a. You can accept the default
 - The default always includes the M.N.P.S.
- b. You can enter the desired `<PBS_EXEC_BASE>`
 - If you enter `<PBS_EXEC_BASE>`, the package manager appends M.N.P.S to your answer.
- c. You can enter the desired `<PBS_EXEC_BASE>/M.N.P.S`.
 - If you enter `<PBS_EXEC_BASE>/M.N.P.S`, the package manager accepts your answer.

2. In the second request, the package manager shows you `<PBS_EXEC_BASE>/M.N.P.S` and asks:

```
You had entered: <PBS_EXEC_BASE>/M.N.P.S
```

```
Enter path to package base directory [?,q]
```

You **must** enter what it shows, which is `<PBS_EXEC_BASE>/M.N.P.S`, exactly.

3.4.7.2.iii Reverting to a Prior Version on Solaris

In order to revert to a prior version of PBS, you must uninstall the new package and reinstall the prior package.

3.4.8 Installing on an Altix Running SuSE

3.4.8.1 Before Installing

Please do not jump straight to this section in your reading. Before downloading and installing PBS, please make sure that you have read the following and taken any required steps:

- Prerequisites: All of Section 2.1, "Prerequisites for Running PBS", Section 2.2, "Single Execution System", Section 2.3, "Multiple Execution Systems", and Section 2.4, "User

Requirements on UNIX/Linux" and their subsections.

- Pre-installation steps: All of Section 3.1, "Requirements for Installing PBS Professional", Section 3.2, "Licensing for PBS Professional", and Section 3.3, "Overview of Installing PBS", and their subsections.
- Installation steps: Section 3.4.1, "Create Required Accounts", Section 3.4.2, "Permissions", Section 3.4.3, "Installation Programs", Section 3.4.4, "If You Have Multiple Execution Hosts", Section 3.4.6, "Installation Procedure (Most Systems)", and Section 3.4.5, "Systems With Prior Versions of PBS", and their subsections.

3.4.8.2 Download and Install the New PBS

1. Download the PBS `tar.gz` package to `/tmp`.

2. Change directory to `/tmp`:

```
cd /tmp
```

3. Extract from the package:

```
tar zxvf PBSPro_<M.N.P.S>-SGIAltix_pp<ProPack version>_ia64.tar.gz
```

4. Change directories:

```
cd PBSPro_<M.N.P.S>
```

5. Execute installation script:

```
./INSTALL
```

6. Answer the questions; supply the license string; do not start the daemons yet.

7. Change directories:

```
cd /opt/pbs/M.N.P.S/sbin
```

8. Rename the standard PBS MOM:

```
mv pbs_mom pbs_mom.bak
```

9. Copy the cpuset PBS MOM to `pbs_mom`:

```
cp -rp pbs_mom.cpuset pbs_mom
```

10. Check to see that new paths work:

```
/opt/pbs/default/sbin/pbs_mom -s list
```

11. Check to see that symbolic links to old paths have been created:

```
<old path>/sbin/pbs_mom -s list
```

3.4.8.3 Start and Test the New PBS

12. Start PBS. If the PBS startup script is not used on the Altix, `pbs_mom` will not start:

```
/etc/init.d/pbs start
```

13. Check to see that vnode definitions for `pbs_mom` have been generated:

```
/opt/pbs/M.N.P.S/sbin/pbs_mom -s list
```

14. Check to see that the PBS daemons are running. You should see that there are three daemons running: `pbs_mom`, `pbs_server`, `pbs_sched`:

```
ps -ef | grep pbs
```

15. Submit jobs as a normal user.

Submit a job to the default queue:

```
echo "sleep 60" | /opt/pbs/default/bin/qsub
```

16. Verify that the jobs are running:

```
/opt/pbs/default/bin/qstat -an
```

3.4.9 Installing PBS on the SGI ICE/XE

3.4.9.1 Before Installing

Please do not jump straight to this section in your reading. Before downloading and installing PBS, please make sure that you have read the following and taken any required steps:

- Prerequisites: All of Section 2.1, "Prerequisites for Running PBS", Section 2.2, "Single Execution System", Section 2.3, "Multiple Execution Systems", and Section 2.4, "User Requirements on UNIX/Linux" and their subsections.
- Pre-installation steps: All of Section 3.1, "Requirements for Installing PBS Professional", Section 3.2, "Licensing for PBS Professional", and Section 3.3, "Overview of Installing PBS", and their subsections.
- Installation steps: Section 3.4.1, "Create Required Accounts", Section 3.4.2, "Permissions", Section 3.4.3, "Installation Programs", Section 3.4.4, "If You Have Multiple Execution Hosts", Section 3.4.6, "Installation Procedure (Most Systems)", and Section 3.4.5, "Systems With Prior Versions of PBS", and their subsections.

3.4.9.2 Considerations

Most sites that are running PBS Professional on an SGI ICE/XE will not want PBS to manage the cpusets on the machine. These sites should install the standard PBS MOM.

Placement sets improve job placement on execution nodes. See section 18.8, "Placement Sets and Task Placement" on page 935 in the PBS Professional Administrator's Guide. If you run the cpuset MOM, placement sets are generated automatically. If you run the standard MOM, placement sets are not automatically generated, and you may wish to configure placement sets. See section 3.4.9.9, "Configuring Placement Sets on the SGI ICE/XE", on page 55 for the steps involved.

3.4.9.3 SGI ICE Components

An ICE system consists of one Admin node, one or more Service (login) nodes, and a set of one or more compute racks. Each compute rack consists of one or more IRU nodes and one or more compute nodes per IRU. The racks are diskless. The root file system of the IRU and compute nodes are mounted read-only from a NAS managed by the Admin node. There is a single image of the root file system for all of the compute nodes and a separate image for all of the IRU nodes. Tempo node management commands are used to publish the image to the various nodes in a process that involves powering down the nodes, pushing a new image, and re-powering the nodes.

In a typical configuration, user home file systems are mounted from NAS, and each node has a separately mounted file system for `/var/spool`.

SGI follows a naming convention when preparing a system for shipment. Service nodes are named "service0", "service1", ... Compute nodes are named `rRiLnN` where 'R' is the rack number starting with 1; 'L' is the IRU node number within a rack starting with 0 in each rack; N is the node number, starting with 0, under the specific Rack Leader. For example, two racks with 2 IRUs per rack and 4 nodes per IRU are named:

Table 3-4: Node Names

	Rack 1	Rack 2
IRU 0	<i>r1i0n0</i>	<i>r2i0n0</i>
	<i>r1i0n1</i>	<i>r2i0n1</i>
	<i>r1i0n2</i>	<i>r2i0n2</i>
	<i>r1i0n3</i>	<i>r2i0n3</i>

Table 3-4: Node Names

	Rack 1	Rack 2
IRU 1	<i>r1i1n0</i>	<i>r2i1n0</i>
	<i>r1i1n1</i>	<i>r2i1n1</i>
	<i>r1i1n2</i>	<i>r2i1n2</i>
	<i>r1i1n3</i>	<i>r2i1n3</i>

3.4.9.4 Requirements for the SGI ICE with ProPack 6 or 7 or Performance Suite 1

- In order to run PBS on the SGI ICE with ProPack 6 or 7 or Performance Suite 1, SGI's Tempo node management tools must already be installed. You will be using the following Tempo commands:

Table 3-5: Tempo Commands

Tempo command	Description
<code>cnodes --compute</code>	List the compute node names; useful in scripting operations
<code>cpower --down NODE</code>	Powers down
<code>cpower --up NODE</code>	Powers up the named nodes
<code>cimage --...</code>	Manages the file system image for the various nodes

- You must use the correct names for the Admin and Service nodes in any commands.
- If you will use PBS to manage the cpusets on the ICE/XE, the file `/etc/sgi-compute-node-release` must be present in order for the PBS `init.d` script to create placement set data and for the `pbs_habitat` script to set up the CPUSET flags.

3.4.9.5 Choosing Whether PBS Will Manage cpusets With SGI ICE Running ProPack 6 or 7 or Performance Suite 1

You can choose whether or not to use PBS to manage the cpusets on the SGI ICE. Most sites will not want PBS to manage the cpusets. To have PBS manage the cpusets, install `pbs_mom.cpuset`; otherwise, install `pbs_mom.standard`.

3.4.9.5.i Using PBS to Manage cpusets on the SGI ICE

If you are running the cpuset MOM, the `init.d/pbs` script will configure one vnode per MOM. This enables cpusets and sets sharing to `default_shared`.

To provide the maximum number of available CPUs on a small node, make sure that the file `/etc/sgi-compute-node-release` is present. This way, on installation the `pbs_habitat` script will add a `“cpuset_create_flags 0”` to Mom's config file.

In order to exclude CPU 0, change the MOM configuration file line to

```
cpuset_create_flags CPuset_CPU_EXCLUSIVE
```

This flag controls only whether CPU 0 is included in the PBS cpuset.

There is only one logical memory pool available per node on the SGI ICE. If, at startup, MOM finds:

- any CPU in an existing, non-root, non-PBS cpuset
- CPU 0 has been excluded as above

MOM will

- Exclude that CPU from the top set `/dev/cpuset/PBSPro`
- Create the top set with `mem_exclusive` set to `false`

Otherwise, the top set is created using all CPUs and with `mem_exclusive` set to `True`.

3.4.9.6 Installation of the PBS Server and Scheduler

The PBS Server, Scheduler and commands are installed on a single service node; here we assume this node is `“service0”`. To install the complete PBS package:

1. Log on to `service0` as root.
2. Unzip and untar the install package.
3. Change to the directory created when the package was untarred.
4. Run the `INSTALL` script: accept the default locations, and select the number `1` (Server, execution and commands) option.
5. When asked if you wish to start PBS, answer `“no”`.
6. Modify `/etc/pbs.conf` by changing the line `“PBS_START_MOM=1”` to `“PBS_START_MOM=0”`.

3.4.9.7 Installation of the PBS MOM

You install and configure MOM on each of the compute nodes. The installation and configuration is only done once on the root file system; this is then pushed to all of the compute nodes.

1. Log on to the Admin node as root.
2. Determine which image file is being used on the compute nodes. To list the nodes on the rack 1:

```
cimage --list-nodes r1
```

It will show output in the form “*node: image_name kernel*” similar to

```
r1i0n0: compute-sles10sp1 2.6.26.46-0.12-smp
```

Thus node r1i0n0 is running the image “*compute-sles10sp1*” and the kernel version “*2.6.26.46-0.12-smp*”. For the remaining steps, it is assumed that those are the images and kernel available.

3. List the available images:

```
cimage --list-images
```

which will list the images available for the compute nodes. Each image may have multiple kernels.

4. Unless you are experienced in managing the image files, it is suggested that you create a copy of the image in use and install PBS in that copy. To copy an image:

```
cinstallman --create-image --clone --source compute-sles10sp1 --image compute-sles10sp1pbs
```

5. The image file lives in the directory `/var/lib/systemimager/images`, so change into the `tmp` directory found in the new image just cloned:

```
cd /var/lib/systemimager/images/  
compute-sles10sp1pbs/tmp
```

6. Copy the PBS install package into this `tmp` directory, unzip and untar it.
7. Change back to the root home directory and `chroot` to the new image file:

```
chroot /var/lib/systemimager/images/  
compute-sles10sp1pbs /bin/sh
```

The new root is in effect.

8. Change directory into the PBS directory which was created below `tmp` when the PBS package was untarred. This is `/tmp/PBSPro_<version>`, which was created for the

image when the PBS package was untarred on the Admin node:

```
cd /tmp/PBSPro_<version>
```

9. Run the INSTALL script to install the PBS programs into the normal execution directory /opt/pbs/M.N.P.S in this system image.

```
./INSTALL
```

- a. Accept the default directories.
- b. Choose install option 2 (execution).
- c. Answer “No” when you are asked whether you wish to start PBS now.

10. The default MOM does not use cpusets. To use cpusets, replace the default pbs_mom executable with PBS_EXEC/sbin/pbs_mom.cpuset.

```
cp PBS_EXEC/sbin/pbs_mom.cpuset PBS_EXEC/sbin/pbs_mom
```

PBS_EXEC is the path to the location where the PBS binaries were installed.

11. Exit from the chroot shell and return to root's normal home directory.

12. Power down each rack of compute nodes:

```
for n in `cnodes --compute` ; do
    cpower --down $n
done
```

13. Publish the new system image to the compute nodes:

```
cimage --push-rack compute-sles10sp1pbs r\*
```

This instruction will take several minutes to finish.

14. Set the new image and kernel to be booted. This set need not be done if: (1) rather than cloning a new image, you have installed PBS into the image already running on the compute nodes; or (2) you are using an image that was already pushed to the nodes.

```
cimage --set compute-sles10sp1pbs 2.6.26.46-0.12-smp r\*i\*n\*
```

15. Power up the compute nodes:

```
for n in `cnodes --compute` ; do
    cpower --up $n
done
```

It will take several minutes for the compute nodes to reboot.

3.4.9.8 Adding Compute Nodes

1. Log on to the service node as root
2. On the Service node, start the PBS Server and Scheduler:

```
/etc/init.d/pbs start
```

3. Using `qmgr`, add the compute nodes to the PBS configuration:

```
for N in `cnodes --compute`  
do  
    qmgr -c "create node $N"  
done
```

3.4.9.9 Configuring Placement Sets on the SGI ICE/XE

If you are not already using PBS to manage the placement sets on the ICE/XE (you are running the standard MOM), you may wish to configure placement sets. This will improve job placement on execution nodes. See section 18.8, "Placement Sets and Task Placement" on page 935 in the PBS Professional Administrator's Guide.

Placement sets can be defined only after you have defined the compute nodes as in the previous section.

1. Shut down the server.
2. Add a resource named "router" to the `PBS_HOME/server_priv/resourcedef` file, by adding the following line:

```
router type=string_array flag=h
```

3. Restart the server
4. Run the placement set generation script:

```
PBS_EXEC/lib/init.d/sgiICEplacement.sh
```

5. Verify the result:
 - a. Run the `pbsnodes -a` command
 - b. Look for the line "`resources_available.router`" in each node. The value assigned to the "router" resource should be in the form "`r#,r##`", where *r* identifies the rack number and *i* identifies the IRU number.

If you are using PBS to manage the cpusets on the ICE/XE (you are running the cpuset MOM), the above does not need to be performed, because the generation of a vnode description automatically provides placement set information.

3.4.10 Installing PBS on the Cray XT

3.4.10.1 Before Installing

Please do not jump straight to this section in your reading. Before downloading and installing PBS, please make sure that you have read the following and taken any required steps:

- Prerequisites: All of Section 2.1, "Prerequisites for Running PBS", Section 2.2, "Single Execution System", Section 2.3, "Multiple Execution Systems", and Section 2.4, "User Requirements on UNIX/Linux" and their subsections.
- Pre-installation steps: All of Section 3.1, "Requirements for Installing PBS Professional", Section 3.2, "Licensing for PBS Professional", and Section 3.3, "Overview of Installing PBS", and their subsections.
- Installation steps: Section 3.4.1, "Create Required Accounts", Section 3.4.2, "Permissions", Section 3.4.3, "Installation Programs", Section 3.4.4, "If You Have Multiple Execution Hosts", Section 3.4.6, "Installation Procedure (Most Systems)", and Section 3.4.5, "Systems With Prior Versions of PBS", and their subsections.

3.4.10.2 About Installation on the Cray XT

This installation procedure differs considerably from the procedures for prior versions of PBS. Please review this section carefully before starting to install PBS Professional on the XT. The installation instructions given in this section are for an initial installation only. If there is an older version of PBS Professional installed on your system, perform an upgrade, following the instructions in section 5.11.1, "Installing PBS To Upgrade on the Cray XT", on page 134.

The PBS Professional package supports CNL.

The PBS Professional Scheduler and Server can be run on an external host; they do not have to run on your Cray system. In the following instructions, the examples show the PBS Server and Scheduler running on the `sd` node.

One instance of the `pbs_mom` daemon is run on each login node of the Cray system, allowing PBS Professional to select the login node to which a job is assigned.

The PBS Professional installation script, `INSTALL`, will detect that it is running on a Cray XT and perform special actions to manage `PBS_EXEC`, `/etc/pbs.conf`, and `PBS_HOME`.

On Cray XT systems only, if the directory `/opt/modulefiles` exists, a module description file is installed in `/opt/modulefiles/pbs/M.N.P.S`.

For information on the functionality provided by modulefiles see the man pages for `module(1)` and `modulefile(4)` on Cray XT systems.

3.4.10.3 Caveats

The PBS Professional software may have been installed on your Cray System as part of the initial setup of the machine. If it is already installed and you perform an initial installation, that may destroy the pre-installed system software.

To verify whether PBS Professional software is currently installed, log in as root to your Cray system and execute the following command:

```
/opt/pbs/default/bin/qstat --version
```

If PBS Professional is already installed, you must follow the steps for an upgrade. See section 5.11.1, “Installing PBS To Upgrade on the Cray XT”, on page 134.

3.4.10.3.i Not Supported

Kerberos is not supported on the Cray XT.

3.4.10.4 Prerequisites

3.4.10.4.i Required Libraries

The XML library, `libexpat.a`, must be available on the login nodes. This is fairly standard in Linux distributions.

3.4.10.4.ii Use Correct Version Number

You must know the version number of the PBS Professional package you are installing. The following examples use version number 10.0.0.08470. The version number of the package you are installing will be different. Please be sure to use the version number of your package.

3.4.10.4.iii Hardware and Software Requirements

Confirm the following before beginning installation:

- You must have root permission.
- All configuration files (such as `/etc/nsswitch.conf`) must have been updated appropriately for your site. The `/etc/nsswitch.conf` file must contain the following lines:

```
hosts: files dns
networks: files dns
```

If the `/etc/nsswitch.conf` file is not updated, the installation will fail and you will receive the following error message:

```
pbsd_main: unable to get my hostname
```

- For SLES 9 based systems, the `nfslock` and `rstatd` services must be enabled at boot

time. This is not required for SLES 10 and beyond. The current SLES release is documented in the `/etc/SuSE-release` file on Cray service nodes. If you are running a SLES 9 based system, verify that `nfslock` and `rstatd` are on:

```
# ssh root@boot
boot::~~# xtopview
default/:// # chkconfig
rstatd    on
nfslock  on
default/:// # exit
```

- The PBS Server and Scheduler must be installed on a node with network access to the Altair License Server. This will allow the PBS Server to use the PBS Professional licenses. This may be the `sdb` node on the XT or it may be a node separate from the XT system.
- On the login nodes and on the node on which the PBS Server and Scheduler will run (typically the `sdb` node), make sure that the `/var` directory is owned by root.
- Determine how you will start up and shut down PBS Professional. You must choose whether to configure resiliency communication agent (RCA) services or use `init.d` (`chkconfig`) to start up PBS Professional.
- Determine whether you need to set up a License Server or request a new PBS Professional license key from Altair Engineering. For more information on generating and installing license keys, see the Altair License Management System Installation and Operations Guide, available on the Altair website at www.altair.com.
- Compute nodes must be in the “batch role” before they can be used by jobs under the control of PBS Professional. To set all of the compute nodes in the batch role:

On the boot node execute:

```
xtprocadm -k m batch
```

Restart ALPS:

```
/etc/init.d/alps restart
```

3.4.10.5 Installing PBS on the Cray XT

To install PBS Professional, follow the steps listed below:

1. Obtain the current distribution of PBS Professional. This is a gzipped tar file, for example:

```
PBSPro_10.0.0-linux24_x86_64_crayxt.tar.gz
```

2. Create a temporary install directory on the boot node:

```
smw:~# ssh root@boot mkdir /rr/current/software/pbs-install
```

3. Copy the distribution file to that directory:

```
smw:~# scp PBSPro_10.0.0-linux24_x86_64_crayxt.tar.gz root@boot:/rr/current/software/pbs-install
```

4. Log in as root to the boot node.

5. Change to the temporary install directory.

6. Unzip and untar the distribution:

```
smw:~# ssh root@boot boot001:~
# cd /rr/current/software/pbs-install
# gunzip PBSPro_10.0.0-linux24_x86_64_crayxt.tar.gz
# tar xf PBSPro_10.0.0-linux24_x86_64_crayxt.tar
```

7. Change into the directory created when the package was untarred:

```
cd PBSPro_10.0.0
```

8. Run the INSTALL script in this directory:

```
# ./INSTALL
```

9. Accept the default location for the installed binaries, for example:

```
PBS_EXEC=PBSPro_10.0.0.08470
```

10. Accept the default location for the PBS working directories , in

```
PBS_HOME=/var/spool/PBS
```

11. When asked for the name of the node on which the server and scheduler will be run, enter the name of the node.

12. When asked for the type of PBS Installation, choose “1” :

```
1. Server, execution and commands.
```

13. When asked for the “License File Location(s)”, follow the instructions in section 3.3.4,

“Setting PBS License Server Location”, on page 28.

14. The post-install procedure is automatically run. You will be asked “Would you wish to start PBS now?” To defer starting PBS, answer “*no*”.

15. Verify that the new PBS Professional configuration file `/etc/pbs.conf` exists in the shared view of the default shared root file system. The file should look like the following, except that “sdb” will be replaced with the PBS Server node you specified:

```
PBS_SERVER=sdb
PBS_EXEC=/opt/pbs/default
PBS_HOME=/var/spool/PBS
PBS_START_SERVER=0
PBS_START_SCHED=0
PBS_START_MOM=0
```

16. Remove the temporary install directory:

```
cd $HOME
rm -fr /rr/current/software/pbs-install
```

The PBS Professional software is now installed, and you are ready to proceed with configuration.

3.4.10.6 Configuring PBS Professional for the Cray XT

The PBS Professional software must be installed before you configure it. The next steps guide you through the configuration process for PBS Professional.

3.4.10.6.i Caveats

`PATH` must be included in the `pbs_environment` file. The `PATH` value is passed on to batch jobs. To maintain security, it is important that `PATH` be restricted to known, safe directories. Do not include “.” in `PATH`.

3.4.10.6.ii Prerequisites to Configuration

If your site has disabled the use of remote operation functions (“r” commands) and output cannot be returned for jobs running on compute nodes, enable the use of the `cp` command by adding `$usecp` to the `$PBS_HOME/mom_priv/config` file on each login node.

3.4.10.6.iii Steps in Configuring PBS for the Cray XT

1. As root, log in to the boot node

```
# ssh root@boot
```

2. Tailor the install for the login nodes, enabling `pbs_mom`:

Use the `xtopview` command:

```
xtopview -m "pbs.conf for login" -c login
```

Run the `xtspec` command to specialize the change to `/etc/pbs.conf` for the login nodes:

```
class/login/: # xtspec /etc/pbs.conf
```

Edit `/etc/pbs.conf` ensuring the following lines are present:

```
PBS_START_MOM=1
```

Make sure that the Server and Scheduler do not start:

```
PBS_START_SERVER=0
```

```
PBS_START_SCHED=0
```

Ensure that the `PBS_SERVER` entry names the node on which you plan to run the Server and Scheduler, and that `PBS_EXEC` is set to `/opt/pbs/default`.

Ensure that `/opt/pbs/default` is a symbolic link to the location you specified for the binaries.

3. Exit from `xtopview`.
4. Determine the node ID , or NID, of the node on which you plan to run the PBS Server and Scheduler. This is usually the `sdb` node. Typically the NID of the `sdb` node is 3.

```
grep sdb /etc/hosts
```

5. Tailor the install for the `sdb` node, enabling `pbs_server` and `pbs_sched`:

Use the `xtopview` command to run the `pbs_postinstall` script:

```
# xtopview -m "pbs.conf for Server" -n <NID>
```

Run the `xtspec` command to specialize the change to `/etc/pbs.conf` for the `sdb` node:

```
<nodeID>: # xtspec /etc/pbs.conf
```

Edit the `/etc/pbs.conf` file or verify that it contains the following lines:

```
PBS_START_SERVER=1
```

```
PBS_START_SCHEDULER=1
```

Verify that the MOM start line in `/etc/pbs.conf` is:

```
PBS_START_MOM=0
```

Ensure that the `PBS_SERVER` entry in `pbs.conf` names the node on which you plan to run the Server and Scheduler and that `PBS_EXEC` is set to `/opt/pbs/default`.

6. Exit from `xtopview`.
7. Log in as root to the node on which the PBS Server `pbs_server` will be running. This is shown here as the `sdb` node.

```
# ssh root@sdb
```

8. At this point, you **must** start the PBS Server with the `init.d/pbs` script. Do not start the server by any other means.

```
sdb # /etc/init.d/pbs start
```

This will create or update the PBS working directories, `PBS_HOME`, in `/var/spool`.

9. Using the `qmgr` command, create the PBS execution node or nodes, which are the login nodes, and set the special XT resource types for those nodes.

When creating a vnode to represent a login node, use the short name returned by the `hostname` command on the login node. For example, if `hostname` returns `HostA`, do the following:

```
Qmgr: create node HostA
```

If you create a vnode with a different name from the short name returned by `hostname`, the following happens:

- MOM creates a vnode whose name is the short name returned by `hostname`
- The vnode you created is not recognized by MOM, and is marked *stale*

Example of creating a vnode to represent a login node:

```
sdb: # module load pbs
```

```
sdb: # qmgr
```

```
Qmgr: create node login1
```

The execution queue `workq` is automatically created during the normal PBS post-install procedure.

10. The default for `kill_delay` is 10 seconds. It is not recommended that you set it to less than 10 seconds.

```
sdb: # qmgr -c "set queue workq kill_delay=<delay>"
```

11. Make sure that `flatuid` is set to true in the Server configuration:

```
sdb: # qmgr -c "list server"
```

If `flatuid` is not set to true, set it to true:

```
sdb: # qmgr -c "set server flatuid=true"
```

12. To ensure that scheduling of jobs is initiated when PBS is restarted, set the Server's scheduling attribute to true:

```
sdb: # qmgr -c "set server scheduling=true"
```

13. Shut down the PBS Services using the following command:

```
/etc/init.d/pbs stop
```

3.4.10.6.iv Configuration on Login Nodes

On each login node, log in and perform the following steps:

1. Start the PBS MOM daemon with the `init.d/pbs` script. Do not use any other means to start `pbs_mom`. This will create or update the PBS working directories, `PBS_HOME`, in `/var/spool`.

```
Login:~ # /etc/init.d/pbs start
```

2. Stop the MOM daemon by running the `init.d/pbs` script:

```
Login:~ # /etc/init.d/pbs stop
```

3. Add the following directives in the `pbs_mom` configuration file, `PBS_HOME/mom_priv/config`:

For CNL:

```
$clienthost sdb
$alps_client /usr/bin/apbasil
```

The sample configuration is now complete.

3.4.10.7 Configuring PBS Professional Startup and Shutdown

To use `init.d` and the `chkconfig` command to start PBS Professional, execute Startup Procedure 1. To use RCA services to start PBS Professional execute Startup Procedure 2.

3.4.10.7.i Startup Procedure 1: Using `init.d` (`chkconfig`) to Start PBS Professional

The PBS script is located in `/etc/init.d/pbs`. To use the `init.d` script and the `chkconfig` command to start up PBS Professional, execute the following steps:

1. Find the node ID (NID) of the node on which the PBS Server and Scheduler will run. If this is the `sdb` node, do this:

```
grep sdb /etc/hosts
```

2. Enable PBS startup for the Server and Scheduler on the `sdb` node. The NID of the `sdb` node is usually 3. This example uses an NID of 3.

```
boot001: # xtopview -n 3
node/3:/ # chkconfig pbs on
node/3:/ # exit
```

3. Enable PBS startup for the MOM on each login node:

```
# xtopview -c login
class/login/: # chkconfig pbs on
class/login/: # exit
```

3.4.10.7.ii Startup Procedure 2: Using RCA to Start PBS Professional

To configure RCA services to start PBS Professional, execute the following procedure:

1. The set of required files and directories that exist in the directory specified by `PBS_HOME` must be created or updated. `PBS_HOME` is normally `/var/spool/PBS`, but it may have been changed during the installation.

Run `{PBS_EXEC}/etc/pbs_habitat` for the Server/Scheduler node, which is usually the `sdb` node.

Run `{PBS_EXEC}/etc/pbs_habitat` for the login nodes.

2. Configure `pbs_mom` to start on each login node. Use the `xtservconfig` command:

```
sdb:~# /opt/xt-boot/default/bin/snos64/xtservconfig -c login add PBS-MOM
```

3. Configure `pbs_server` and `pbs_sched` to start on the appropriate node. For example, the following will run them on the `sdb` node (NID 3). See the `xtservconfig(8)` man page for more information.

```
sdb:~# /opt/xt-boot/default/bin/snos64/xtservconfig -n 3 add PBS-SERV
sdb:~# /opt/xt-boot/default/bin/snos64/xtservconfig -n 3 add PBS-SCHED
```

4. The scheduling process starts automatically when PBS Professional is started. To prevent this and have the scheduling start when the login nodes are booted, include the `-a`

false option to the SVC_START_CMD line in the /opt/xt-boot/default/etc/pbs-server startup script. For example:

```
SVC_START_CMD="/opt/pbs/default/sbin/pbs_server -a false"
```

3.4.10.8 Configuring PBS to Support Cray

PBS provides tight integration with Cray. To use this integration, configure PBS to support the Cray. For your next step, go to section 11.5, "Support for Cray Systems" on page 666 in the PBS Professional Administrator's Guide.

3.5 Installation on Windows Systems

3.5.1 Before Installing

Please do not jump straight to this section in your reading. Before downloading and installing PBS, please make sure that you have read the following and taken any required steps:

- Prerequisites: All of Section 2.1, "Prerequisites for Running PBS", Section 2.2, "Single Execution System", Section 2.3, "Multiple Execution Systems", Section 2.5, "Recommended PBS Configurations for Windows", Section 2.6, "Windows User Authorization", and Section 2.7, "Windows Caveats" and their subsections.
- Pre-installation steps: All of Section 3.1, "Requirements for Installing PBS Professional", Section 3.2, "Licensing for PBS Professional", and Section 3.3, "Overview of Installing PBS", and their subsections.

3.5.2 Installation Locations

On 32-bit Windows systems, PBS is installed in `\Program Files\PBS Pro\`, and on 64-bit Windows systems, PBS is installed in `\Program Files (x86)\PBS Pro\`. We use `INSTALLATION_BASE` to indicate `\Program Files\` or `\Program Files (x86)\`. When reading these instructions, replace `INSTALLATION_BASE` with the appropriate directory.

When PBS is installed on a complex, the MOM must be run on each execution host. The Server and Scheduler only need to be installed on one of the hosts or on a front-end system. For Windows clusters, PBS is provided in a single package containing the following:

- PBS Professional software
- Supporting text files (software license, README, release notes, etc.)

3.5.3 PBS Windows Considerations

PBS Professional is supported if the domain controller server is configured “native”. Running PBS in an environment where the domain controllers are configured in “mixed-mode” is not supported.

You must install PBS Professional from an Administrator account.

Before you install PBS on Windows, make sure you are using the correct type of account. See section 2.5.2, “Windows Configuration in a Domained Environment”, on page 14.

PBS Professional requires that the drive that PBS was installed under (e.g. `\Program Files\PBS Pro` or `\Program Files (x86)\PBS Pro`) be configured as an NTFS filesystem.

Before installing PBS Professional, be sure to uninstall any old PBS Professional files. For uninstalling versions 5.4.2 through 8.0, use a domain admin account. For details see “Uninstalling PBS Professional on Windows” on page 75.

You can specify the destination folder for PBS using the “Ask Destination Path” dialog during setup. After installation, icons for the `xpbs` and `xpbsmon` GUIs will be placed on the desktop and a program file menu entry for PBS Professional will be added. You can use the GUIs to operate on PBS or use the command line interface via the command prompt.

The PBS data service account is the same as the PBS service account. You do not need to create a separate data service account.

3.5.4 Pre-installation Configuration

Before installing PBS Professional on a Windows cluster, perform the following system configuration steps first.

The following discussion assumes that the `pbs_server` and `pbs_sched` services will be installed on a front-end host called “hostA”, and the `pbs_mom` service will be installed on all the vnodes in the complex that will be running jobs, “hostB ... hostZ”.

1. Be sure that hostA, hostB, ..., hostZ consistently resolve to the correct IP addresses. A wrong IP address to hostname translation can cause errors for PBS. Make sure the following are done:
 - Configure your system to talk to a properly configured and functioning DNS server
 - Add the correct host entries to the following files:

On Windows XP:

```
c:\windows\system32\drivers\etc\hosts
```

For example, if your Server is `serverA.example.com` with address 192.0.0.231, then add the entry:

192.0.0.231 serverA.example.com serverA

2. Set up any user accounts that will be used to run PBS jobs. They should not be Administrator-type of accounts, that is, not a member of the “Administrators” group so that basic authentication using `hosts.equiv` can be used.

The accounts can be set up using either of the following:

Start->Control Panel->Administrative Tools->Computer Management->Local Users & Groups

- or -

Start->Control Panel->User Manager

3. Once the accounts have been set up, edit the `hosts.equiv` file on all the hosts to include hostA, hostB, ..., hostZ to allow accounts on these hosts to access PBS services, such as job submission and remote file copying.

The `hosts.equiv` file can usually be found in either of the following locations:

C:\winnt\system32\drivers\etc\hosts.equiv

C:\windows\system32\drivers\etc\hosts.equiv

4. Create required accounts. Before you install PBS, you must create the accounts that PBS requires. PBS has different requirements for these accounts, depending on whether it is in a domained or standalone environment.

For installation in a domained environment, see section 2.5.2.2, “Installation Account”, on page 15 and section 2.5.2.3, “The PBS Service Account”, on page 15.

For installation in a standalone environment, see section 2.5.3.2, “Installation Account”, on page 18 and section 2.5.3.3, “The PBS service account for Standalone Environments”, on page 18.

Create the following accounts, following the instructions for the environment:

installation account

The installation account is the account from which you will execute the PBS `INSTALL` program.

PBS service account

The PBS service account will actually run the PBS services: `pbs_server`, `pbs_mom`, `pbs_sched`, and `pbs_rshd`. The PBS service account is also recommended for performing any updates of the PBS configuration files.

3.5.5 Software Installation

Next you will need to install the PBS software on each execution host of your complex. The PBS Professional installation program will walk you through the installation process.

3.5.5.1 Installation Account

PBS must be installed from a local or domain account. A local account has local administrator privilege on the install host (i.e. the account must be a member of the local "Administrators" group on the local machine). A domain account has either domain administrator privileges (i.e. a member of "Domain Admins" group) or is a member of the local Administrators group on the local (install) host. The `INSTALL` program must be executed from a domain account which is a member of the "Domain Admins". See section 2.5, "Recommended PBS Configurations for Windows", on page 13.

3.5.5.2 Installation Steps

1. If you are installing from the PBS Professional CD-ROM, insert the CD-ROM into your computer's CD-ROM drive, browse to your CD-ROM drive, and click on the PBS Professional program icon.

Alternatively, you can download the latest PBS Professional package from the PBS Web site, and save it to your hard drive. Run the self-extracting `pbspro.exe` package, and then the installation program, as shown below.

```
Admin> PBSPro_10.0.0-windows.exe
```

On Windows XP, Service Pack 2 (SP2), upon launching the installer, a window may be displayed saying the program is from an unknown publisher. In order to proceed with the installation of PBS, click the “*Run*” button.

2. Review and accept the License Agreement, then click *Next*.
3. Supply your Customer Information, then click *Next*.
4. Review the installation destination location. You may change it to a location of your choice, provided the new location meets the requirements stipulated in section 2.5, “Recommended PBS Configurations for Windows”, on page 13. Then click *Next*.
5. When installing on an execution host in the complex, select the “Execution” option from the *install* tool, then click *Next*.
6. You will be prompted for the name of the PBS service account. Enter your choice, then click *Next*.
7. You will then be prompted to enter a password for the PBS service account (as discussed in the previous section). The password typed will be masked with “*”. An empty pass-

word will not be accepted. Enter your chosen password twice as prompted, then click *Next*.

You may receive the following “error 2245” when PBS creates the PBS service account. This means “The password does not meet the password policy requirements. Check the minimum password length, password complexity and password history requirements.”

You *must* use the same password when installing PBS on additional execution hosts as well as on the PBS Server host.

8. The installation tool will show two screens with informative messages. Read them; click *Next* on both.
9. On the “Editing PBS.CONF file” screen, specify the hostname on which the PBS Server service will run, then click *Next*.
10. On the “Editing HOSTS.EQUIV file” screen, follow the directions on the screen to enter any hosts and/or users that will need access to this local host. Then click *Next*.
11. On the “Editing PBS MOM config file” screen, follow the directions on the screen to enter any required MOM configuration entries (as discussed in “MOM Parameters” on page 263 of the PBS Professional Reference Guide). Then click *Next*.
12. Lastly, when prompted, select *Yes* to restart the computer and click *Finish*.

Repeat the above steps for each execution host in your complex. When complete you are ready to install PBS Professional on the host that will become the PBS Server host.

1. Install PBS Professional on hostA, selecting the “*All*” option. Next, you will be prompted for your license server location. Following this, the install program will prompt for information needed in setting up the `hosts.equiv` file, etc. Enter the information requested for hosts hostB, hostC, ..., hostZ, clicking *Next* to move between the different input screens.
2. Finally, run `pbsnodes -a` on hostA to see if it can communicate with the execution hosts in your complex. If some of the hosts are seen to be down, then go to the problem host and restart the MOM, using the commands:

```
Admin> net stop pbs_mom
Admin> net start pbs_mom
```

3.5.6 Post Installation Considerations

The installation process will automatically create the following file,

```
[PBS Destination folder]\pbs.conf
```

containing at least the following entries:

```
PBS_EXEC=[PBS Destination Folder]\exec
PBS_HOME=[PBS Destination Folder]\home
PBS_SERVER=server-name
```

where `PBS_EXEC` will contain subdirectories where the executable and scripts reside, `PBS_HOME` will house the log files, job files, and other processing files, and `server-name` will reference the system running the PBS Server. The `pbs.conf` file can be edited by calling the PBS program “`pbs-config-add`”. For example, on 32-bit Windows systems:

```
\Program Files\PBS Pro\exec\bin\pbs-config-add "PBS_SCP=\winnt\scp.exe"
```

or on 64-bit Windows systems:

```
\Program Files (x86)\PBS Pro\exec\bin\pbs-config-add  
"PBS_SCP=\winnt\scp.exe"
```

Don't edit `pbs.conf` directly as the permission on the file could get reset causing other users to have a problem running PBS.

The auto-startup of the services is controlled by the PBS `pbs.conf` file as well as the Services dialog. This dialog can be invoked via selecting *Settings->Control Panel->Administrative Tools->Services*. If the services fail to start up with the message, “incorrect environment”, it means that the `PBS_START_SERVER`, `PBS_START_MOM`, and `PBS_START_SCHED` `pbs.conf` variables are set to *0 (False)*.

Upon installation, special files in PBS home directory are set up so that some directories and files are restricted in access. The following directories will have files that will be readable by the `\Everyone` group but writable only by Administrators-type accounts:

```
PBS_HOME/server_name
PBS_HOME/mom_logs/
PBS_HOME/sched_logs/
PBS_HOME/spool/
PBS_HOME/server_priv/accounting/
```

The following directories will have files that are only accessible to Administrators-type of accounts:

```
PBS_HOME/server_priv/
PBS_HOME/mom_priv/
PBS_HOME/sched_priv/
```

You should review the recommended steps for setting up user accounts and home directories, as documented in section 2.6, “Windows User Authorization”, on page 20.

3.5.7 Windows XP SP2 Firewall

Under Windows XP Service Pack 2 (SP2) the Windows Firewall may have been turned on by default. If so, it will block incoming network connections to all services including PBS. Therefore after installing PBS Professional, to allow `pbs_server`, `pbs_mom`, `pbs_sched`, and `pbs_rshd` to accept incoming connections:

Access *Settings->Control Panel->Security Center->Windows Firewall*, and verify that the Windows Firewall has been set to “*ON*” to block incoming network connections.

From this panel, you can either turn Windows Firewall “*off*”, or click on the *Exceptions* tab and add the following to the list:

```
[INSTALL PATH]\exec\sbin\pbs_server.exe
[INSTALL PATH]\exec\sbin\pbs_mom.exe
[INSTALL PATH]\exec\sbin\pbs_sched.exe
[INSTALL PATH]\exec\sbin\pbs_rshd.exe
```

3.5.8 Windows Services

The Windows version of PBS contains a service called `pbs_rshd` for supporting remote file copy requests for delivering job output and error files to destination hosts. (Keep in mind that `pbs_rshd` does not allow normal `rsh` activities but only `rcp`.) This service is used only by PBS; it is not intended to be used by administrators or users.

`pbs_rshd` will read either the `%WINDIR%\system32\drivers\etc\hosts.equiv` file or the user's `.rhosts` file for determining the list of accounts that are allowed access to the localhost during remote file copying. PBS uses this same mechanism for determining whether a remote user is allowed to submit jobs to the local Server. `pbs_rshd` is started automatically during installation but can also be started manually by typing either of the following two commands:

```
net start pbs_rshd
```

-or-

```
pbs_rshd -d
```

This latter form of invocation runs `pbs_rshd` in debug mode where logging output will be displayed on the command line.

If user1 on *hostA* causes a file to be copied to *hostB* (running `pbs_rshd`) so that the source is `file1` and the destination is `hostB:file2`, `pbs_rshd` authenticates user1 as follows:

- If *user1* is a non-administrator account (e.g. not belonging to the Administrators group),

then the copy will succeed in one of 2 ways:

- *user1@hostA* is authenticated via *hostB*'s `hosts.equiv` file; or
- *user1@hostA* is authenticated via user's `[PROFILE_PATH]/.rhosts` on *hostB*.

See also section 2.6.2, "Windows User's HOMEDIR", on page 22.

One of the following must be configured:

- The system-wide `hosts.equiv` file on *hostB* includes as one of its entries:
`hostA`
- or, `[PROFILE_PATH]\.rhosts` on *userA*'s account on *hostB* includes:
`hostA userA`
- If *userA* is an administrator account, or the source is `file1` and the target is `userB@hostB:file2`
then use of the account's `[PROFILE_PATH]\.rhosts` file is the only way to authenticate, and it needs to have the entry:
`hostA userA`

These two methods of authentication are further discussed in the **PBS Professional User's Guide**.

3.5.9 Network Drives and File Delivery

If users require jobs to have output or error files going into some network location, and that network location is mapped to the same local drive (for instance drive `Q`), then you need to put the following two lines in MOM's config file. (For additional information on MOM configuration parameters, see "Contents of MOM Configuration File" on page 264 of the PBS Professional Reference Guide.)

```
$usecp *:Q: Q:
$usecp *:q: q:
```

The above causes any job output or error files having the form, "`<hostname>:<the letter "q">:file-path`" to be passed to the copy command as:

```
Q:file-path
```

or

```
q:file-path
```

The reason for putting a wildcard entry for *hostname* in `$usecp` is to get around the possibility of MOM seeing different permutations of *hostname* for the destination host. The upper and lower cases of "q" are needed in order to get a match in all possible situations.

3.5.10 Changing the PBS service account Password

Normally, the PBS service account password must not be changed. But if it is deemed necessary to change it perhaps due to a security breach, then do so using the following steps:

First, change the PBS service account's password on a machine in a command prompt from an admin-type of account by typing:

Domain environments:

```
net user <name of PBS service account> * /domain
```

Non-domain environment:

```
net user <name of PBS service account> *
```

Then the Service Control Manager (SCM) must be provided with the new password specified above. This can be done via the GUI-based Services application found as one of the Administrative Tools, or unregister and re-register the PBS services with password.

On 32-bit Windows systems:

```
pbs_account --unreg "\Program Files\PBS Pro\exec\sbin\pbs_server.exe"
pbs_account --unreg "\Program Files\PBS Pro\exec\sbin\pbs_mom.exe"
pbs_account --unreg "\Program Files\PBS Pro\exec\sbin\pbs_sched.exe"
pbs_account --unreg "\Program Files\PBS Pro\exec\sbin\pbs_rshd.exe"
```

```
pbs_account --reg "\Program Files\PBS Pro\exec\sbin\pbs_server.exe"
pbs_account --reg "\Program Files\PBS Pro\exec\sbin\pbs_mom.exe"
pbs_account --reg "\Program Files\PBS Pro\exec\sbin\pbs_sched.exe"
pbs_account --reg "\Program Files\PBS Pro\exec\sbin\pbs_rshd.exe"
```

On 64-bit Windows systems:

```
pbs_account --unreg "\Program Files (x86)\PBS
Pro\exec\sbin\pbs_server.exe"
pbs_account --unreg "\Program Files (x86)\PBS Pro\exec\sbin\pbs_mom.exe"
pbs_account --unreg "\Program Files (x86)\PBS Pro\exec\sbin\pbs_sched.exe"
pbs_account --unreg "\Program Files (x86)\PBS Pro\exec\sbin\pbs_rshd.exe"
```

```
pbs_account --reg "\Program Files (x86)\PBS Pro\exec\sbin\pbs_server.exe"
pbs_account --reg "\Program Files (x86)\PBS Pro\exec\sbin\pbs_mom.exe"
pbs_account --reg "\Program Files (x86)\PBS Pro\exec\sbin\pbs_sched.exe"
pbs_account --reg "\Program Files (x86)\PBS Pro\exec\sbin\pbs_rshd.exe"
```

The register form (last four lines above) can take an additional argument `-p password` so that you can specify the password on the command line directly.

3.5.11 Uninstalling PBS Professional on Windows

For uninstalling the old PBS with version between 5.4.2 and 8.0, an account that is a member of the “Domain Admins” group will still have to be used.

To remove PBS from a Windows system, do one of the following:

- Go to *Start Menu->Control Panel->Add/Remove Programs* (Windows XP) menu and select the PBS Professional entry and click “*Change/Remove*”
- Double-click on the PBS Windows installation package icon to execute. This will automatically delete any previous installation.

If the uninstallation process complains about not completely removing the PBS installation directory, then remove it manually, for example by typing:

```
cd \Program Files
```

or

```
cd \Program Files (x86)
rmdir /s "PBS Pro"
```

Under some conditions, if PBS is uninstalled by accessing the menu options (discussed above), the following error may occur:

```
...Ctor.dll: The specified module could not be found
```

To remedy this, do the uninstall by running the original PBS Windows installation executable (e.g. `PBSPro_10.0.0-windows.exe`), which will remove any existing instance of PBS.

During uninstallation, PBS will not delete the PBS service account because there may be other PBS installations on other hosts that could be depending on this account. However, the account can be deleted manually using an administrator-type of account as follows:

In a domain environment:

```
net user <name of PBS service account> /delete /domain
```

In a non-domain environment:

```
net user <name of PBS service account> /delete
```

At the end of uninstallation, it is recommended to check that the PBS services have been completely removed from the system. This can be done by opening up the Services dialog. In Windows XP, go to *Start Menu->Control Panel->Performance and Maintenance->*

Administrative Tools->Services and check to make sure `PBS_SERVER`, `PBS_MOM`, `PBS_SCHED`, and `PBS_RSHD` entries are completely gone. If any one of them has a state of "*DISABLED*", then you must restart the system to get the service removed.

During uninstallation, you can choose to remove all files in `PBS_HOME/datastore`.

3.6 Enabling Metascheduling with HPC Basic Profile

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

See section 7.1, "Definitions" on page 481 in the PBS Professional Administrator's Guide for definitions of HPCBP elements.

3.6.1 Multiple HPC Basic Profile Servers

You must create one HPCBP MOM for each HPCBP Server that the PBS complex will use. You will not be able to enable the HPCBP MOM if jobs are running on the HPCBP MOM host. To prevent PBS from scheduling jobs on this host, you can set state of the vnode to *OFFLINE* before starting the MOM, then clear the *OFFLINE* state once the HPCBP MOM is enabled.

To create an HPCBP MOM, preventing PBS from running jobs on the host:

- Install PBS on the HPCBP MOM host, choosing the execution host (MOM) option. When asked if you want to start PBS, do not do so.
- Using the `qmgr` command, create the HPCBP MOM vnode. See section 3.7.1, "Creating or Modifying Vnodes", on page 78.
- Using the `pbsnodes -o` command, set the vnode running the HPCBP MOM to *OFFLINE*.
- Configure the MOM to be an HPCBP MOM. This is done by configuring four vnode attributes on the vnode where the HPCBP MOM will run. Instructions for configuring an HPCBP MOM are given in section 7.3, "Configuring PBS for HPCBP" on page 486 in the PBS Professional Administrator's Guide.
- Start the HPCBP MOM using the PBS start script. See section 6.1, "Starting and Stopping PBS: UNIX and Linux", on page 167.
- Using the `pbsnodes -r` command, clear the *OFFLINE* state of the HPCBP MOM.

3.6.2 Platforms for PBS Hosts When Using HPCBP

If you wish to use HPCBP with PBS, you must use supported platforms for the following:

- The PBS Server
- Any machine that will host an HPCBP MOM
- Any machine from which HPCBP jobs will be submitted

See the Release Notes for a list of supported architectures.

3.6.2.1 HPCBP Server Platform

The HPCBP Server can be any HPC Basic Profile compliant server.

3.6.3 OpenSSL Requirement

You must install elements of OpenSSL in order to be able to use metascheduling with HPC Basic Profile. The following table lists which machines require these elements, and what the elements are:

Table 3-6: Machines Requiring OpenSSL

Machine	Required OpenSSL Installation
HPCBP MOM host	OpenSSL runtime libraries OpenSSL package containing the <code>c_rehash</code> utility
PBS Server host	OpenSSL runtime libraries
PBS submission hosts used to submit HPCBP jobs	OpenSSL runtime libraries
PBS hosts used to run <code>pbs_password</code> or <code>pbs_migrate_users</code> for HPCBP users	OpenSSL runtime libraries

3.7 After Installation

3.7.1 Creating or Modifying Vnodes

After `pbs_server` is started, the vnode list may be created via the `qmgr` command. First start up `pbs_mom`, then use `qmgr` to add the vnode. For example, to add a new vnode, use the “create” sub-command of `qmgr`:

```
Qmgr: create node vnode_name [attribute=value]
```

where the attributes and their possible values are listed in “Vnode Attributes” on page 360 of the PBS Professional Reference Guide. Vnode attributes cannot be used as vnode names. On a multi-vnode system, only the natural vnode should be created this way.

All comma-separated attribute-value strings must be enclosed in quotes.

Below are several examples of creating vnodes via `qmgr`.

```
Qmgr: create node mars resources_available.ncpus=2  
Qmgr: create node venus
```

3.7.1.1 Modifying Vnodes

Once a vnode has been created, its attributes and/or boolean resources can be modified using the following `qmgr` syntax:

```
Qmgr: set node vnode_name [attribute[+|-]=value]
```

where attributes are the same as for `create`. For example:

```
Qmgr: set node mars resources_available.inner=true  
Qmgr: set node mars resources_available.haslife =true
```

3.7.1.2 Deleting Vnodes

Nodes can be deleted via `qmgr`, using the `delete node` syntax, as the following example shows:

```
Qmgr: delete node mars  
Qmgr: delete node pluto
```

3.7.1.3 Caveats

When setting a vnode’s attributes, follow the rules given in section 3.5, “How to Configure MOMs and Vnodes”, on page 45 of the PBS Professional Administrator’s Guide.

It is not a good idea to try to use `qmgr` to create the vnodes for an Altix, other than the natural vnode. You do need to create the natural vnode via `qmgr`. It is possible to use `qmgr` to create a vnode with any name. The "[x]" naming does not imply any special significance; it just an internal convention for naming vnodes on an Altix. The fact that you can create a vnode with a weird name does not mean however that the MOM on the host knows about that vnode. If the MOM does not know about the vnode, the vnode will be considered "stale" and not usable. By default, MOM only knows about the natural vnode, the one whose name is the same as the host.

3.7.2 Post-Installation Validation

If you wish to validate the installation of PBS Professional, at any time, run the `pbs_probe` command. It will review the installation (installed files, directory and file permissions, etc) and report any problems found. For details, see “`pbs_probe`” on page 74 of the PBS Professional Reference Guide. (The `pbs_probe` command is not available under Windows.)

Use the `qstat` command to find out what version of PBS Professional you have.

```
qstat -fB
```

3.7.3 Command and Man Page Paths

Set paths for all users to include PBS commands and man pages. For paths, see section 3.3.8, “Default Installation Locations”, on page 32. Administrator commands are in `PBS_HOME/sbin`, and user commands are in `PBS_HOME/bin`. Man pages are `PBS_HOME/man`.

3.7.4 Supporting X Forwarding for Interactive Jobs

To support X forwarding, edit each MOM’s `PATH` variable to include the directory containing the `xauth` utility. X forwarding is not available under Windows.

3.7.5 Troubleshooting: Server Becomes Unresponsive on Startup

If the server becomes unresponsive on startup, the server may be trying to contact the wrong license server. For recovery instructions, see section 14.1, “Server Host Bogs Down After Startup” on page 757 in the PBS Professional Administrator’s Guide.

Chapter 4

Licensing

As of 11.0, PBS jobs are licensed using a new Altair license server, and you **must** install and configure this new license server for PBS to function. This licensing is independent of application licensing. Applications are licensed by whichever server is required by the application.

4.1 Introduction

4.1.1 Overview of Licensing for PBS Jobs

The Altair license server makes a number of floating licenses available to run PBS jobs on hosts in a network. PBS uses a units-based or token-based licensing scheme. The number of licenses needed is proportional to the number of CPUs requested by a job. All PBS licenses are now floating licenses, and do not depend on which host is being used. Licenses are used for the CPUs used by a job, instead of for hosts. It is no longer the hosts that are licensed, it is the CPUs used by a job.

Multiple PBS complexes can use the same Altair license server.

4.1.2 Tracking License Use

The PBS server can keep licenses available locally. It checks the licenses out of the Altair license server. The licenses are checked out of the server, whether or not they are being used by jobs. The PBS server hands them out to jobs as needed. The PBS server keeps track of these licenses in the `license_count` server attribute. In this attribute, the PBS server tracks how many licenses are checked out by the PBS server in the `Avail_Local` parameter, and how many are in use in the `Used` parameter. See “Server Attributes” on page 310 of the PBS Professional Reference Guide.

The minimum number of licenses to keep checked out is specified in the `pbs_license_min` server attribute.

The number of seconds to keep unused CPU licenses, when the number of licenses is above `pbs_license_min`, is specified in `pbs_license_linger_time`.

If licenses are released when a job finishes, the PBS server waits `pbs_license_linger_time` seconds before checking them back into the license server, during which time they will be kept under the `avail_local` parameter of the `license_count` attribute.

The `Avail_Global` parameter of the `license_count` attribute tracks how many licenses are kept by the Altair license server; that is, that have not been checked out from the license server.

4.1.3 Recommendations

We recommend that you set `pbs_license_min` to the number of CPUs in the PBS complex. This way, other applications can't grab the licenses and prevent PBS from running a job.

4.1.4 License Server Versions and PBS Releases

Each feature in the license file has a version number associated with it. You need to have a license file in which the feature version is at least as new as the PBS Professional version. You can use PBS Professional when its version is older than or the same as the features in the license file.

You will need an Altair license file when you upgrade PBS Professional or when your license expires. If you need additional licenses in the mean time, contact Altair Support.

4.1.5 Old Licensing Model Used for Pre-9.0 Versions and Trial Licenses

The old PBS licensing scheme required that a license key be provided to PBS via "`PBS_SERVER/server_priv/license_file`". This key determined whether a job could be run on a number of CPUs on one or more hosts. This proprietary license model accepted either a regular license or a trial license, which had a short expiration time. For regular licenses, there were 3 types: UNIX/Windows license (type '5'), Linux licenses (type 'L'), and floating licenses (type 'F').

The old PBS proprietary license model will continue to license PBS Professional versions 8.0 and earlier. Customers who upgrade from version 8.0 and earlier to 9.0 and later will need to download, install and configure the Altair license server. See section 3.2, "Licensing for PBS Professional", on page 26.

4.2 Definitions

Floating License

A unit of license dynamically allocated (checked out) when a user begins using an application on some host (when the job starts), and deallocated (checked in) when a user finishes using the application (when the job ends). The floating licenses discussed in this chapter license PBS.

License Manager Daemon (**lmx-serv-altair**)

Daemon that functions as the license server. See the Altair License Management System Installation and Operations Guide, available at the Altair website, www.altair.com.

License Server

Manages licenses for PBS jobs.

License Server List Configuration

One form of redundant license server configuration. A collection of license server files, or "<port>@<host>" settings, pointing to license servers managing Altair licenses. Each server on the list is tried in turn. There could be X licenses on <server1>, Y licenses on <server2>, and Z licenses on <server3>, and the total licenses available would actually be X+Y+Z, but a request must be satisfied only by one server at a time. The first running server is the only server queried.

Memory-only Vnode

Represents a node board that has only memory resources (no CPUs), for example, an Altix memory-only blade.

Redundant License Server Configuration

Allows licenses to continue to be available should one or more license servers fail. There are two types: 1) license server list configuration, and 2) three-server configuration.

Three-server Configuration

One form of redundant license server configuration. Means that if any 2 of the 3 license servers are up and running (referred to as a quorum), the system is functional, with 1 server acting as master who can issue licenses. If the master goes down, then another server must take over as master. This is set up as a license file on each of the 3 redundant servers containing:

```
SERVER <server1> ... <port1>
SERVER <server2> ... <port2>
SERVER <server3> ... <port3>
```

PBS Professional can point to a license server host that has these license file entries.

4.3 Licensing Requirements

In order to check out licenses, the PBS server needs the following directories to be writable by root:

`/tmp`

`/var/tmp`

root's home directory

If it is not possible to have root's home directory be writable, the workaround is to create the `.x-formation` directory by hand. For example, on a Cray system:

```
% xtopview
```

```
% cd /root
```

```
% mkdir .x-formation
```

4.4 Configuring PBS for Licensing

4.4.1 Configuring the Server for Licensing

To configure the PBS server for licensing:

1. Specify the license server port and hostname by setting the server's `pbs_license_info` attribute. See “Server Licensing Attributes” on page 85 and “Setting Server Licensing Attributes” on page 87.
2. Optionally, specify the minimum number of licenses to keep permanently checked out of the license server and ready to be used by jobs, by setting the server's `pbs_license_min` attribute. It is recommended that you set this to the total number of CPUs in the complex.

This is the total for all `resources_available.ncpus` configured for each vnode. See section 4.4.3.2, “Setting `pbs_license_min`”, on page 88.

Setting `pbs_license_min` to the total number of CPUs in the complex is especially useful if you are running both PBS and HyperWorks from one license server. This way, other applications won’t prevent PBS from starting jobs.

Also see section 4.7.4, “Licensing and Reservations”, on page 97.

3. Optionally, specify the maximum number of CPUs to be licensed at any time by setting the server’s `pbs_license_max` attribute.
4. Optionally, specify the number of seconds to keep an unused license before returning it to the pool when there are more than the minimum number of licenses checked out by setting the server’s `pbs_license_linger_time` attribute.

4.4.2 Server Licensing Attributes

Licensing is controlled by the server attributes listed below:

FLicenses

The number of floating CPU licenses available to PBS. Equal to the `Avail_Global` + `Avail_Local` of the `license_count` attribute. Integer. Set by the server. Readable by all. Default: zero.

license_count

`license_count= Avail_Global:<X> Avail_Local:<Y> Used:<Z>
High_Use:<W>`

Avail_Global is the number of PBS CPU licenses still kept by the Altair License Server (checked in).

Avail_Local is the number of PBS CPU licenses in the internal PBS license pool (checked out).

Used is the number of PBS CPU licenses currently in use.

High_Use is the highest number of CPU licenses checked out and used at any given time while the current instance of the PBS server is running.

“*Avail_Global*” + “*Avail_Local*” + “*Used*” is the total number of CPU licenses configured for one PBS complex.

Integer. Set by Server. Readable by all. Default: zero.

pbs_license_info

Hostname of license server, or local pathname to the actual license file(s), which is associated with a license server. String. Set by PBS Manager. Readable by all. Default value: empty string, meaning no server to contact.

To set `pbs_license_info` to the hostname of the license server:

```
Qmgr: set server
      pbs_license_info=<port1>@<host1>:<port2>@<host2>:...:<portN>@<
      hostN>
```

where <host1>, <host2>, ..., <hostN> can be IP addresses.

Windows: Use semicolons (;) instead of colons (:), and enclose the path list in double quotes if it contains any spaces or when you are listing more than one port/server. For example:

```
Qmgr: set server
      pbs_license_info="<port1>@<host1>;<port2>@<host2>;...;<portN>@
      <hostN>"
```

To unset `pbs_license_info` value:

```
Qmgr: unset server pbs_license_info
```

`pbs_license_linger_time`

The number of seconds to keep an unused CPU license, when the number of licenses is above the value given by `pbs_license_min`. Time. Set by PBS Manager. Readable by all. Default: 3600 seconds.

To set `pbs_license_linger_time`:

```
Qmgr: set server pbs_license_linger_time=<Z>
```

To unset `pbs_license_linger_time`:

```
Qmgr: unset server pbs_license_linger_time
```

`pbs_license_max`

Maximum number of licenses to be checked out at any time, i.e. maximum number of CPU licenses to keep in the PBS local license pool. Sets a cap on the number of CPUs that can be licensed at one time. Long. Set by PBS Manager. Readable by all. Default: maximum value for an integer.

To set `pbs_license_max`:

```
Qmgr: set server pbs_license_max=<Y>
```

To unset `pbs_license_max`:

```
Qmgr: unset server pbs_license_max
```

`pbs_license_min`

Minimum number of licenses to permanently keep checked out of the license server and ready to be used by jobs, i.e. the minimum number of CPU licenses to keep in the PBS local license pool. It is recommended that you set `pbs_license_min` to the total number of CPUs in your complex.

This is the total for all `resources_available.ncpus` configured for each vnode. Long. Set by PBS Manager. Readable by all. Default: zero.

To set `pbs_license_min`:

```
Qmgr: set server pbs_license_min=<X>
```

To unset `pbs_license_min`:

```
Qmgr: unset server pbs_license_min)
```

4.4.3 Setting Server Licensing Attributes

The administrator can define the following server attributes via `qmgr`:

4.4.3.1 Setting the Licensing Information in `pbs_license_info`

To set `pbs_license_info` to the hostname of the license server(s), run the following as a PBS manager or administrator (i.e. root):

```
Qmgr: set server pbs_license_info=
      <port1>@<host1>:<port2>@<host2>:...:<portN>@ <hostN>
```

where `<host1>`, `<host2>`, ..., `<hostN>` can be IP addresses.

Use the above if you have one license server or are using the license server list form of redundancy.

The default value is empty, meaning no external license server is contacted.

These actions follow:

- If the `pbs_license_info` is currently set to a non-empty value, then it is set to the new value. All previous licenses are checked back into the previous license server, and the connection to that server is terminated.
- If the `pbs_license_info` is currently unset, then it is set to the new value. If the PBS complex was licensed by a trial license key, then the trial licensing is discontinued. An attempt is made to initialize connection to the license server whose location is described in `pbs_license_info`. If the connection fails, `qmgr` and `server_logs` contain a warning:

```
"Unable to connect to license server at pbs_license_info=<X>"
```

The attribute is still set to this new value.

Upon successful connection to the license server, PBS tries to re-license the running jobs (which ran previously via trial licenses). Any jobs that are currently running continue to run, even if not all the necessary licenses are obtained for these jobs.

Windows: Use semicolons (;) instead of colons (:), and enclose the path list in double quotes if it contains any spaces or when you are listing more than one port/server. For example:

```
qmgr> set server pbs_license_info=
    "<port1>@<host1>;<port2>@<host2>;...;<portN>@<hostN>"
```

To unset `pbs_license_info`, simply run the following as PBS Manager or Administrator (i.e. root):

```
Qmgr> unset server pbs_license_info
```

These actions follow:

- The `pbs_license_info` attribute is set to the empty string, previous licenses are checked back into the previous license server, and connection to this license server is shut down.
- Currently running jobs are allowed to run to completion.

Once PBS is installed, the license file location can be changed (if, for example, the Altair license server is moved to another host), by setting the following server attribute via `qmgr`:

```
Qmgr: set server pbs_license_info= <port1>@<host1>
```

4.4.3.2 Setting `pbs_license_min`

It is recommended that you set `pbs_license_min` to the total number of CPUs in your complex. This is the total for all `resources_available.ncpus` configured for each vnode. To set `pbs_license_min`, run the following as a PBS manager or administrator (i.e. root):

```
Qmgr: set server pbs_license_min=<X>
```

These actions follow:

- If `<X>` is not numeric, or is less than 0, or is greater than the value of `pbs_license_max`, then the attribute is not set, and `qmgr` outputs an error and returns a non-zero value.
- The next time PBS updates its license pool (usually every 5 minutes or when a job is started/exited), the new value for `pbs_license_min` is known internally to the PBS server.

If the server cannot obtain the amount of CPU licenses given by `pbs_license_min` from the license server, then it will try to obtain as many as possible, log the error, and keep trying to get more up to the correct value over some period of time.

To unset `pbs_license_min`, run the following as a PBS manager or administrator (i.e. root):

```
Qmgr: unset server pbs_license_min
```

This will cause `pbs_license_min` to revert to its default value.

The next time PBS updates its license pool (usually every 5 minutes or when a job is started/exited), the value of `pbs_license_min` will have a default value, and unused licenses are returned to the license server. The return of the license(s) to the pool is constrained by `pbs_license_linger_time`.

4.4.3.3 Setting `pbs_license_max`

To set `pbs_license_max`, run the following as a PBS manager or administrator (i.e. root):

```
Qmgr: set server pbs_license_max=<Y>
```

These actions follow:

If `<Y>` is not numeric, or is less than 0, or is less than the value of `pbs_license_min`, then the attribute is not set, and `qmgr` outputs an error and returns a non-zero value.

The next time PBS updates its license pool (usually every 5 minutes or when a job is started/exited), the new value for `pbs_license_max` is known internally to the PBS server.

If the new value for `pbs_license_max` is less than the previous value, the value is not adjusted by taking away the licenses that are currently in use by running jobs. This allows the jobs to continue to run. However, as jobs exit, the PBS server will adjust the internal license pool to have no more than the number of licenses given by `pbs_license_max`.

To unset `pbs_license_max`, run the following as a PBS manager or administrator (i.e. root):

```
Qmgr: unset server pbs_license_max
```

This causes the `pbs_license_max` value to revert to its default value. The next time PBS updates its license pool (usually every 5 minutes or when a job is started/exited), the value of `pbs_license_max` is set to its default.

4.4.3.4 Setting `pbs_license_linger_time`

To set `pbs_license_linger_time`, run the following as a PBS manager or administrator (i.e. root):

```
Qmgr> set server pbs_license_linger_time=<Z>
```

These actions follow:

- If `<Z>` is not numeric, or is less than or equal to 0, then the attribute is not set, and `qmgr` outputs an error and returns a non-zero value.
- The next time PBS updates its license pool (usually every 5 minutes or when a job is started/exited), the new value for `pbs_license_linger_time` is known internally to the PBS server.

To unset `pbs_license_linger_time`, run the following as PBS manager or administrator (i.e. root):

```
Qmgr: unset server pbs_license_linger_time
```

This causes the value of `pbs_license_linger_time` to revert to its default. The next time PBS updates its license pool (usually every 5 minutes or when a job is started/exited), the value of `pbs_license_linger_time` is at the default.

4.4.4 Licensing and PBS Server Failover

The server attribute values for `pbs_license_info`, `pbs_license_min`, `pbs_license_max`, and `pbs_license_linger_time` are set through the primary server. Since these values are saved in a shared location, the secondary server can use these licensing parameters. No additional licensing steps are needed for the secondary server to work properly.

4.5 Redundant License Servers

The Altair license server can be part of a redundant license server system. See the Altair License Management System Installation and Operations Guide, available at the Altair website, www.altair.com.

If you have redundant license servers, some or all of the licenses used by your site will be available if a server host crashes. You can use the PBS Server/Scheduler host as one of the license server hosts. There are two ways to set up your license servers for redundancy:

1. The three-server redundant setup
2. The license server list

In the three-server setup, as long as at least two of the servers are up, PBS can get all the licenses. In this arrangement, each of the servers has the same license file. One server acts as “master”.

In the license server list arrangement, each server has some of the licenses, and PBS will try each in turn. The first running server is the only server queried. This arrangement is better where the network is unreliable.

Set up redundancy before starting PBS.

When choosing license server hosts, choose machines that are stable and unlikely to be rebooted, and that have excellent communication.

For information on redundant license servers, see the Altair License Management System Installation and Operations Guide, available on the Altair website at www.altair.com.

4.5.1 To Use the Three-server Redundant Setup

1. Make sure that the three servers are configured correctly according to the guidelines for the three-server arrangement.
2. Make sure that the PBS server has a file with the contents:

```
SERVER    <host1> ...<port1>
SERVER    <host2> ...<port2>
SERVER    <host3> ...<port3>
USE_SERVER
```

This file can have any name and any location as long as it resides on the PBS server. It is recommended that you put it in `$PBS_HOME/server_priv`.

3. Set the server's `pbs_license_info` attribute to point to the file described in Step 2:

```
Qmgr: set server pbs_license_info= <location of three-server host
file>
```

4.5.2 To Use the License Server List

1. Make sure that the servers are configured correctly according to the guidelines for the server list.
2. List the servers in order: *primary:secondary:tertiary*. Set the server's `pbs_license_info` attribute to point to the list:

```
Qmgr: set server pbs_license_info=
<port1>@<host1>:<port2>@<host2>:<port3>@<host3>:...:<portN>
@<hostN>
```

For Windows, use semicolons instead of colons, and enclose the path list in double quotes if any of the paths contains spaces.

4.6 Displaying Licensing Information

4.6.1 Viewing License Information in Server Attributes

To see the information in the server attributes including `Flicenses`, `pbs_license_info`, `pbs_license_min`, etc, run the following:

```
qstat -Bf
```

or

```
qmgr -c "list server"
```

4.6.2 Viewing the Number of Available Licenses

If license tools such as `lmxendutil` are installed on the PBS server host, the user can use them to discover the number of actual licenses available. These are reported in units that are a multiple of the number of CPUs. For information about viewing licenses if `lmxendutil` is not installed on the PBS server host, see the Altair License Management System Installation and Operations Guide, available at the Altair website, www.altair.com.

4.7 PBS Jobs and Licensing

4.7.1 Examples of Licensing PBS Jobs

The following examples show how PBS licenses jobs.

1. Fit in a single host such as an Altix, packed on the fewest vnodes:

```
qsub -l ncpus=10:mem=20gb -l place=pack
```

This job requests 10 CPUs, so PBS will check out 10 licenses.

2. Request 4 chunks, each with 1 CPU and 4 gb of memory taken from anywhere:

```
qsub -l select=4:ncpus=1:mem=4gb -l place=free
```

The job is requesting 4 CPUs, so PBS will check out 4 licenses.

3. Request 4 vnodes where the arch is *linux* and each vnode is on a separate host:

```
qsub -l select=4:mem=2gb:ncpus=2:arch=linux -l place=scatter
```

Each vnode has 2 CPUs and 2GB memory allocated to the job. The job requests 8 CPUs so PBS will check out 8 licenses.

4. Request to run on a specific host:

```
qsub -l select=1:ncpus=2:mem=50gb:host=zooland
```

The job requests 2 CPUs, so PBS needs to check out 2 licenses.

5. Request resources on different hosts:

```
qsub -l select=2:ncpus=3:mem=6gb -l place=scatter
```

The job requests 6 CPUs, so PBS will check out 6 licenses to run the job.

6. Cpuset: An odd-size job that will fit on a single Altix, but not on any one nodeboard, and the request is not shared:

```
qsub -l select=1:ncpus=3:mem=6gb -l place=pack:excl
```

The job is specifically requesting 3 CPUs, so PBS needs to check out 3 licenses.

7. Cpuset: Request a small number of CPUs but a large amount of memory, exclusively:

```
qsub -l select=1:ncpus=1:mem=25gb -l place=pack:excl
```

The job is requesting 1 CPU, so PBS checks out 1 license.

8. Cpuset: Align a large job within one router, if it fits within a router.

```
qsub -l select=1:ncpus=100:mem=200gb -l place=pack:group=router
```

The job requests 100 CPUs, requiring PBS to check out 100 licenses.

9. Assignment of resources involving memory-only vnodes (nodeboard):

Given 3 vnodes available where

V1 has 4 CPUs and 2 GB

V2 has no CPUs and 3 GB

V3 has no CPUs and 4 GB

V4 has 2 CPUs and 5GB

V1, V2, V3, V4 have sharing attribute set to “*default_shared*”

Case 1: Job requesting no CPUs:

```
qsub -l select=1:ncpus=0:mem=3gb
```

This job would not run since it asks for no CPUs. No licenses would be checked out. Note that that even though a chunk within a job can ask for 0 CPUs, the job as a whole must request at least one CPU.

Case 2: Job requesting a combination of vnodes with CPUs and memory-only vnodes, non-exclusively:

```
qsub -l select=1:ncpus=1:mem=7gb
```

Even though this job may get assigned 2 GB of V1, 3 GB of V2, and 2 GB of V3, the job has requested 1 CPU so PBS needs to check out only 1 license.

Case 3: Job requesting a combination of vnodes with CPUs and memory-only vnodes **exclusively**:

```
qsub -l select=1:ncpus=1:mem=14gb -l place=pack:excl
```

Even though the job may get assigned V1, V2, V3, and V4, the job has requested only 1 CPU so PBS needs to check out 1 licenses.

Case 4: Job not requesting memory-only vnodes:

```
qsub -l select=1:ncpus=2 -lplace=pack:excl
```

Even though this job may get assigned all CPUs (4) of V1 exclusively, this job has only requested 2 CPUs so PBS will check out 2 licenses.

Sub-Case 4a:

```
qsub -lselect=1:ncpus=2:mem=10gb -lplace=excl
```

Scheduler runs job on:

V2: 0 CPU 1 GB

V3: 0 CPU 4 GB

V4: 3 CPUs 5 GB

The job has exclusive use of the vnodes, but the user only requested 2 CPUs, so 2 licenses are needed to run the job.

Sub-Case 4b:

```
qsub -l select=1:ncpus=2:mem=10gb -lplace=excl
```

Scheduler runs job on:

V1: 0 CPU 1 GB

V3: 0 CPU 4 GB

V4: 3 CPUs 5 GB

It doesn't matter how the scheduler assigned the vnodes, the number of CPUs requested is always honored for figuring the number of licenses to get. This also requires 2 licenses.

10. Complex request:

```
qsub -l select=2:ncpus=1+5:ncpus=2
```

This requires 12 licenses.

11. Licensing virtual CPUs:

Given hostA which has 2 physical CPUs, with `resources_available.ncpus` set to 4 on the vnode representing the host:

```
qsub -l select=1:ncpus=3:host=hostA
```

PBS will check out 3 licenses for the job.

12. Licensing on a Hyperthreaded Machine:

A hyperthreaded hostB has 2 physical CPUs, where the OS reports 4 logical CPUs. The vnode representing the host has the default value for `resources_available.ncpus` of 2 and `pcpus=2`. Now set `resources_available.ncpus=4`:

```
qsub -l select=ncpus=4:host=hostB
```

PBS will need 4 licenses to run the job.

13.Licensing multi-core systems:

On multi-core systems, the number of CPUs requiring licenses is the number requested by the job.

On a host with two dual core chips (4 CPUs), a job asking for 2 CPUs would require 2 licenses.

4.7.2 Licensing and Job States

When a job finishes execution (i.e. job has exited), the licenses assigned to the job is returned to the PBS license pool. The pool is then subject to the constraints given in `pbs_license_min` and `pbs_license_linger_time`.

A job that is held is treated as if the job has finished execution. Licenses assigned to it are released to the PBS local license pool.

The licenses assigned to a suspended job also return to the PBS license pool in order to be available to other PBS jobs. The returned licenses are also subject to the `pbs_license_min` and `pbs_license_linger_time` constraints.

The scheduler makes sure licenses are available before resuming any job. If the licenses are not available, a scheduler log message and job comment are provided to warn of the situation.

Example:

JobA has been submitted as follows:

```
qsub -l select=4:ncpus=1:mem=2gb <job.script>
```

The PBS server checked out 4 licenses and executed JobA.

At some point, the scheduler decided to suspend JobA to execute higher priority jobs. The 4 licenses checked out initially are checked back into the PBS license pool (not the license server.) They can be reused by other jobs, and also be made available when JobA gets resumed.

JobA successfully resumes after high priority jobs complete, and PBS was able to get 4 licenses from its pool of licenses.

4.7.3 PBS Jobs and HyperWorks

PBS jobs and Hyperworks applications check out licenses from the same server, but the licenses have distinguishing “features”, as in “PBSprofessional” feature GWUs (GridWorks units), and “Hyperworks” feature HWUs (HyperWorks Units).

4.7.4 Licensing and Reservations

For advance and standing reservations to work, set the `pbs_license_min` to the total number of CPUs, including virtual CPUs, in the PBS complex.

When the scheduler confirms a reservation, the server makes sure the resources requested by the reservation (e.g. vnodes/CPUs) be available for the reservation. If the PBS server always has licenses for the total number of CPUs in the complex always checked out, this guarantees that there will be licenses available to satisfy these reservations.

Upon confirming a reservation, the PBS server will give a warning if `pbs_license_min` is less than the total number of CPUs in the complex.

If the `pbs_license_min` attribute has not been set in the manner recommended, then when the reservation starts, there's a chance that some of the reservation jobs won't run due to a shortage of licenses available. In this case, a job comment and scheduler log entry will be provided.

Example:

The user creates the following reservation requesting 10 CPUs:

```
pbs_rsub -R 1500 -E 1600 -lselect=ncpus=10  
<resv-id> confirmed
```

The user submits jobs to the reservation queue:

```
qsub -q <resv_queue> -l select=ncpus=5 compute_job  
qsub -q <resv_queue> -l select=ncpus=5 compute_job2
```

At the start of the reservation, the two jobs run, each getting 5 licenses.

4.7.5 Suspended Jobs

If you set the server's `pbs_license_min` attribute to the number of CPUs in the complex, suspended jobs will be able to resume when they are ready. Jobs are suspended, and release their licenses, for preemption, and via `qsig -s suspend`. See section 5.9.6.2, "Job Suspension and Resource Usage", on page 284 of the PBS Professional Administrator's Guide.

4.8 Upgrading

4.8.1 Overlay Upgrades

For information about licenses and upgrading, see section 5.3, "Licensing", on page 103.

4.9 Troubleshooting Licensing Errors

4.9.1 Server Becomes Unresponsive on Startup

If the server becomes unresponsive on startup, the server may be trying to contact the wrong license server. For recovery instructions, see section 14.1, "Server Host Bogs Down After Startup" on page 757 in the PBS Professional Administrator's Guide.

4.9.2 Licensing and Loss of Communication to License Server

If PBS loses contact with the Altair License Server, any jobs currently running will not be interrupted or killed. The PBS server will continually attempt to reconnect to the license server, and re-license the assigned vnodes once the contact to the license server is restored.

No new jobs will run if PBS server loses contact with the License server.

If PBS cannot detect a license server host and port when it starts up, the server logs an error message:

```
"Did not find a license server host and port (pbs_license_info=<X>). No
  external license server will be contacted"
```

If the PBS scheduler cannot obtain the licenses to run or resume a job, the scheduler will log a message:

```
"Could not run job <job>; unable to obtain <N> CPU licenses. avail
  licenses=<Y>"
"Could not resume <job>; unable to obtain <N> CPU licenses. avail
  licenses=<Y>"
```

If PBS cannot contact the license server, the server will log a message:

```
"Unable to connect to license server at pbs_license_info=<X>"
```

If the value of the `pbs_license_min` attribute is less than the number of CPUs in the PBS complex when a reservation is being confirmed, the server will log a warning:

```
"WARNING: reservation <resid> confirmed, but if reservation starts now, its
  jobs are not guaranteed to run as pbs_license_min=<X> < <Y> (# of CPUs
  in the complex)"
```

If the PBS server cannot get the number of licenses specified in `pbs_license_min` from the license server, the server will log a message:

```
"checked-out only <X> CPU licenses instead of pbs_license_min=<Y> from
  license server at host <H>, port <P>. Will try to get more later."
```

If the PBS server encounters a proprietary license key that is of not type “-T”, then the server will log the following message:

```
"license key #1 is invalid: invalid type or version".
```

4.9.3 Manually Refreshing License Counts

If the license server becomes temporarily unavailable, PBS may not be able to license new jobs. Normally the scheduler won't query the license server for licenses until a new job is submitted or deleted. However, you can force PBS to contact the license server and refresh the license counts by doing the following:

```
qmgr -c "set server pbs_license_file_location = <...>"
```

4.9.4 User Error Messages

If a user's job could not be run due to unavailable licenses, the job will get a comment:

```
"Could not run job <job>; unable to obtain <N> CPU licenses.
  avail_licenses=<Y>"
```

4.9.5 Checking Trial License Lifetime

If you are not sure how many more days your trial license will last, do the following:

1. Using `qmgr`, clear the comment:

```
qmgr -c "unset server comment"
```
2. Restart the Server
3. Recheck the message in the server's comment attribute.

The comment now contains a message with the number of days left in the trial license. The trial (T) license is good for a maximum of 90 days from the installation of PBS Professional, not from when the license was installed.

Chapter 5

Upgrading PBS Professional

This chapter shows how to upgrade from a previous version of PBS Professional. If PBS Professional is not installed on your system, you can skip this chapter.

5.1 Types of Upgrades

There are two types of upgrades available for PBS Professional. In both, the installation program creates a versioned directory and puts the new PBS binaries in that directory. The existing PBS binaries can remain where they are, except on Solaris, where the package manager allows only one version at a time to be installed.

overlay upgrade

Installs the new `PBS_HOME` on top of the old one. The new `PBS_EXEC` is in its own versioned directory. Jobs stay in place, and can continue to run, except on the Altix.

migration upgrade

Installs the new `PBS_HOME` in a separate location from the old `PBS_HOME`. The new `PBS_HOME` can be in the standard location if the old version has been moved. Jobs are moved from the old server to the new one, and cannot be running during the move. Must be used for Windows.

For specific upgrade recommendations and updates, see the Release Notes.

5.1.1 Systems Requiring Migration Upgrade

Usually, UNIX systems can have overlay upgrades. Some systems and circumstances require migration upgrades. The following require migration upgrades:

- The Cray
- Windows machines
- When moving between 32-bit and 64-bit versions of PBS

5.1.1.1 Upgrading on the Cray

When upgrading PBS on a Cray from any prior version of PBS to 11.0 release of PBS, a migration upgrade is required. An overlay upgrade will not convert mpp* jobs, and existing queued jobs will not execute properly.

In a migration upgrade, it is possible (but unlikely) for jobs that were previously queued to be rejected by the new PBS installation during the `qmove` process. If this happens, check the old PBS system for any jobs that were not successfully moved to the new PBS installation. A job can fail to move to the new system for the following reasons:

- The job uses a mixture of `-lmpp*` and `-lselect/-lplace` syntax
- The `mppnodes` range is ordered from high to low (i.e. `mppnodes='8-6'`)
- There are no `vnodes` whose `PBScraynid` values match the given `mppnodes` values

If it is not clear why a job failed to move, look in the new system's server logs for more details.

5.1.1.2 Upgrading on Windows

For instructions on upgrading on Windows, see section 5.12, “Upgrading Under Windows”, on page 145.

5.2 Differences from Previous Versions

5.2.1 Change in Licensing

As of 11.0, PBS is licensed using a **new** Altair license server. This license server **must** be installed and configured before installing PBS. See section 3.2, “Licensing for PBS Professional”, on page 26 and Chapter 4, “Licensing”, on page 81.

5.2.2 Location of `PBS_EXEC` and `PBS_HOME`

`PBS_EXEC` is the directory that contains the PBS binaries. The PBS installation program creates a versioned directory for the new version of `PBS_EXEC` and puts the binaries there. `PBS_HOME` is the directory where PBS information is stored. The location of `PBS_EXEC` has changed. Please read section 3.3.3, “Versioned Installation”, on page 28, section 3.3.7, “Terminology and Conventions”, on page 32, and section 3.3.8, “Default Installation Locations”, on page 32. Also please read section 3.4.5.2, “Automatic Archiving and User Path Preservation”, on page 37 and section 3.4.5.1.i, “Switching To New Version Now”, on page 36.

5.2.3 Conversion of Properties

The server will convert the old style properties, used in PBS Professional 7.0 and before, to boolean resources.

5.2.4 Change in Home Directory for AIX

The default `PBS_HOME` directory for AIX was changed from

```
/usr/spool/PBS
```

to

```
/usr/local/spool/PBS.
```

5.2.5 Change in Server Start During Overlay Upgrade

During an overlay upgrade, the PBS server must be started so that the server is initialized correctly. If you start PBS using the start script, this is handled automatically. The instructions in this manual for overlay upgrading specify using the start script. However, if, for some reason, you start the server using the `pbs_server` command when overlay upgrading from a pre-11.0 version to 11.0 or later, you must start the server using the `-t updatedb` option to `pbs_server`. In addition, if the automatic update process fails, you must manually start `pbs_server` with the `-t updatedb` option.

5.3 Licensing

Before you perform the upgrade, the server’s `pbs_license_info` attribute may be unset or set to some previous value. You **must** set it to the correct value. See section 4.4.3, “Setting Server Licensing Attributes”, on page 87.

You **must** install, configure, and start the new Altair license server before starting PBS, so follow the steps for installing and starting the license server before upgrading. This license server is new for 11.0. Do not attempt to use an old license server.

5.3.1 Overlay Upgrades

On an overlay upgrade, the old MOM is killed, and the new MOM is started with the "pbs_mom -p" option to allow existing jobs to run. When the 9.0 or later server comes up, the new server will also allow already running jobs to continue. PBS will try to re-license the jobs, but won't kill the jobs if there are not enough licenses.

5.4 Caveats and Advice

5.4.1 Making Time to Upgrade

If you want to avoid having to work around running jobs when you perform an upgrade, you can set PBS up so that there are no running jobs when you want to do the upgrade. Follow these steps:

1. Figure out what the longest-running jobs are likely to need for **walltime**, e.g. two weeks
 2. Pick a time greater than that in the future, e.g. 3 weeks
 3. On all PBS hosts, create dedicated time or a reservation for the amount of time you think the upgrade will require, e.g. a day
- You can use a dedicated time slot, making it so that no jobs will be scheduled for that dedicated time. The system can be shut down all at once at the start of the dedicated time. See section 18.12, "Defining Dedicated Time", on page 974 of the PBS Professional Administrator's Guide.
 - You can create a reservation that reserves an entire host by using `-l place=exclhost`. The following reservation creates a reservation for the host mars, from 10am to 10pm:

```
pbs_rsub -R 1000 -D 12:00:00 -l select = host=mars -l place=exclhost
```

For more on creating reservations, see "pbs_rsub" on page 79 of the PBS Professional Reference Guide.

5.4.2 Multi-vnode Jobs Won't Survive Overlay Upgrade

Multi-vnode jobs cannot survive an overlay upgrade. A multi-vnode job is killed and requeued when any of its MOMs is restarted. If the MOM on the primary vnode is restarted, the job is killed and requeued according to the `node_fail_requeue` server attribute; see section 10.4.2, "Node Fail Requeue: Jobs on Failed Vnodes" on page 636 in the PBS Professional Administrator's Guide.

To avoid this, you can allow the jobs to finish before upgrading. You can set vnodes offline, and then wait until their jobs finish, then upgrade.

5.4.3 Nodes File and Downgrading

If you downgrade after upgrading to version 11 and later, you must use the nodes file from the older version of PBS. The version 11 and later nodes file is incompatible with version 10.4 and earlier servers.

5.4.4 Upgrading Hooks

If you are performing an overlay upgrade, you do not need to take any steps for hooks. Your hooks will be left untouched.

If you are performing a migration upgrade, you must print the hooks out before the upgrade, then read them in after the upgrade.

```
qmgr -c 'print hook' > hook.qmgr
qmgr < hook.qmgr
```

5.4.5 License Information

Before you perform the upgrade, the server's `pbs_license_info` attribute may be unset or set to some previous value. You **must** set it to the correct value. See section 4.4.3, "Setting Server Licensing Attributes", on page 87. You must install and start the new Altair license server before starting PBS, so follow the steps for installing and starting the license server before upgrading. This license server is new for 11.0. Do not attempt to use an old license server.

5.4.6 Upgrading All MOMs Together

Starting with version 9.0, the PBS server and all MOMs must be upgraded at the same time.

5.4.7 Do Not Unset `default_chunk.ncpus`

Do not unset the value for the `default_chunk.ncpus` server attribute. It is set by the server to 1. You can set it to another non-zero value, but a value of 0 will produce undefined behavior. When the PBS Server initializes and the `default_chunk` server attribute has not been specified during a prior run, the Server will internally set the following:

```
default_chunk.ncpus=1
```

This ensures that each "chunk" of a job's select specification requests at least one CPU.

If the Administrator explicitly sets the `default_chunk` server attribute, that setting will be retained across server restarts.

5.5 Before Upgrading from 10.x or Older to 11.x or Later

5.5.1 Create Data Service Account for UNIX/Linux

When upgrading from 10.x or older to 11.x or later, create an account for the data service. For UNIX/Linux, follow the guidelines in section 3.4.1, "Create Required Accounts", on page 34. On Windows, the data service account is the same as the PBS service account.

5.5.2 Get New License File

When upgrading from 10.x or older to 11.x or later, you must have a new license file for the new license server. Contact Altair support.

5.5.3 Run New License Server

When upgrading from 10.x or older to 11.x or later, you **must** install, configure, and start the new Altair license server before starting PBS, so follow the steps for installing and starting the license server before upgrading. This license server is new for 11.0. Do not attempt to use an old license server.

5.5.4 Set Server's `pbs_license_info` Attribute

When upgrading from 10.x or older to 11.x or later, before you perform the upgrade, the server's `pbs_license_info` attribute may be unset or set to some previous value. You **must** set it to the correct value. See section 4.4.3.1, “Setting the Licensing Information in `pbs_license_info`”, on page 87.

5.6 After Upgrading

If you were using `vmem` at the queue or server level before the upgrade, then after upgrading you must add `vmem` to the new `resource_unset_infinite sched_config` option. Otherwise jobs requesting `vmem` will not run.

You may wish to make the upgrade transparent for users, if the installation program hasn't done that already. See section 3.4.5.2.i, “Archiving Prior Version”, on page 37 and section 3.4.5.2.iii, “Re-creating User Paths By Hand”, on page 38.

5.7 Introduction to Upgrading Under UNIX/Linux

When you get your new version of PBS, unpack it (`unzip`, `untar`) as a non-privileged user. When you follow the upgrading instructions below, all of the steps should be performed as root.

5.7.1 Directories

The locations of `PBS_HOME` and `PBS_EXEC` are specified in the file `/etc/pbs.conf`. In the following instructions, replace `PBS_HOME` or `PBS_EXEC` with the appropriate values.

For example, if `pbs.conf` specifies `PBS_HOME` as `/var/spool/PBS`, and an instruction says

```
"mv PBS_HOME PBS_HOME.old",
```

then type

```
"mv /var/spool/PBS /var/spool/PBS.old".
```

5.7.2 Managing Integrations

If you used the `pbsrun_wrap` mechanism with your old version of PBS, you must first unwrap any MPIs that you wrapped. This would include MPICH-GM, MPICH-MX, Intel MPI, MPICH2, etc.

If you are upgrading from a version prior to 9.0 and you had linked the `poë` command to `pbs_poe`, you must remove the link. The new version of PBS uses the `pbsrun_wrap` method; see the PBS Professional Administrator's Guide.

You can wrap your MPIs after upgrading PBS. See the PBS Professional Administrator's Guide.

5.7.3 Upgrading on Multiple Machines

Instead of running the installer by hand on each machine, you can use a command such as `pdsh` with NFS-mounted common directories on all hosts to distribute the installation to each host. The general form to distribute commands to a large number of hosts is

```
# for hosts node001-005
pdsh -w node[001-005] command
```

When using the PBS `INSTALL` command you can make a simple text file called “`answers`” with the answers to the installation prompts, and redirect them into the `INSTALL` script:

```
./INSTALL < answers
```

where `answers` contains the answers for a MOM-only installation:

```
<newline>
<newline>
2
y
<server name>
y
y or n (use appropriate answer)
```

The complete command line would be:

```
pdsh -w node[001-005] ./INSTALL < answers
```

5.7.4 Upgrading on an Altix

When upgrading on an Altix that will be using cpusets, follow the instructions in section 5.9.1, “Upgrading on an Altix or a Complex Containing One or More Altixes”, on page 113. When upgrading on an Altix that will not be using cpusets, follow the instructions in section 5.10, “Migration Upgrade Under UNIX and Linux”, on page 119.

5.8 Overlay Upgrade Under UNIX/Linux

Except on the Altix, use the following steps to perform an overlay upgrade. For the Altix, see section 5.9.1, “Upgrading on an Altix or a Complex Containing One or More Altixes”, on page 113. You will probably want to keep any running jobs running. Multi-vnode jobs cannot survive an overlay upgrade. A multi-vnode job is killed and requeued when any of its MOMs is restarted.

The following commands must be run as “root”.

5.8.1 Back Up Your Existing PBS

Make a tarfile of the PBS_HOME and PBS_EXEC directories.

1. Make a backup directory:

```
mkdir /tmp/pbs_backup
```

2. Make a tarfile of PBS_HOME:

```
cd PBS_HOME/..  
tar -cvf /tmp/pbs_backup/PBS_HOME_backup.tar PBS_HOME
```

3. Make a tarfile of PBS_EXEC:

```
cd PBS_EXEC/..  
tar -cvf /tmp/pbs_backup/PBS_EXEC_backup.tar PBS_EXEC
```

4. Make a copy of your configuration file. If you have configured failover, do this for each server:

```
cp /etc/pbs.conf /tmp/pbs_backup/pbs.conf.backup
```

5. If they exist (PBS 8.0 and later), make a copy of the Version 2 configuration files:

```
mkdir /tmp/pbs_backup/mom_configs  
$PBS_EXEC/sbin/pbs_mom -s list | egrep -v '^PBS' | while read file  
do  
    $PBS_EXEC/sbin/pbs_mom -s show $file > /tmp/pbs_backup/mom_configs/  
    $file  
done  
PBS_EXEC
```

6. Make a copy of the scheduler's directory to modify:

```
cp -r PBS_HOME/sched_priv /tmp/pbs_backup/sched_priv.work
```

5.8.2 Shut Down Your Existing PBS

Shut down PBS, keeping running jobs running. The `qterm` command will use the `-t quick` option unless you specify otherwise.

(PBS versions PBSPro_5.4.0 and later)

```
qterm -m -s -f
```

(PBS versions prior to PBSPro_5.4.0)

```
qterm -m -s
```

If you wish to requeue and/or kill running jobs during shutdown, see section 6.1.7, “Impact of Shutdown / Restart on Running UNIX/Linux Jobs”, on page 176.

5.8.3 Install the New Version of PBS

5.8.3.1 Cray

For Cray, follow the instructions in section 5.11.1, “Installing PBS To Upgrade on the Cray XT”, on page 134.

5.8.3.2 Solaris

Solaris requires a slightly different set of steps when you run the installation program.

1. See section 3.4.7.2, “Notes About Solaris”, on page 46
2. From the directory containing the installation script:

```
./INSTALL
```

The installation program will pick up your existing installation parameters from `/etc/pbs.conf`, and prompt you to confirm that you wish to use them.

When the install script prompts you whether or not to start PBS, if you are on a system running a MOM, answer “*n*” for “no”. Do **not** start `pbs_mom` now. Instead use `pbs_mom -p`, which will preserve running jobs. See the `pbs_mom(8B)` man page.

5.8.3.3 Other UNIX/Linux Platforms

For other UNIX/Linux platforms, install the new version of PBS on all hosts without uninstalling the previous version. The installation program will read your existing installation parameters from `/etc/pbs.conf`, and prompt you to confirm that you wish to use them.

On each host, go to the directory where you put the PBS installation script. Type:

```
./INSTALL
```

When the install script prompts you whether or not to start PBS, if you are on a system running a MOM, answer “*n*” for “no”. Do **not** start `pbs_mom` now. Instead use `pbs_mom -p`, which will preserve running jobs. See the `pbs_mom(8B)` man page.

5.8.4 Prepare the New Scheduler's Configuration File

1. Make a copy of the new `sched_config`, which is in `PBS_EXEC/etc/pbs_sched_config`.
cp PBS_EXEC/etc/pbs_sched_config PBS_EXEC/etc/pbs_sched_config.new
2. Update `PBS_EXEC/etc/pbs_sched_config.new` with any modifications that were made to the current `PBS_HOME/sched_priv/sched_config`. This is saved in the backup directory `/tmp/pbs_backup/sched_priv/work`
3. If it exists, replace the `strict_fifo` option with `strict_ordering`. If you do not, a warning will be printed in the log when the scheduler starts.
4. If you copied over your scheduler log filter setting, make sure the new configuration file has 1024 added to it. If the value is less than 1024, add 1024 to it:

Edit `PBS_EXEC/etc/pbs_sched_config.new`.

If your previous log filter line was:

```
log_filter:256
```

change it to:

```
log_filter:1280
```

If you do not, you will be inundated with logging messages.

5. Move `PBS_EXEC/etc/pbs_sched_config.new` to the correct name and location, i.e. `PBS_HOME/sched_priv/sched_config`.

```
mv PBS_EXEC/etc/pbs_sched_config.new PBS_HOME/sched_priv/sched_config
```

5.8.5 Modify the New Server's Resource File

Add the "h" flag to those vnode-level resources listed in the server's `PBS_HOME/server_priv/resourcedef` file that have the "n" or "f" flag. For example, if you had:

```
switch type=string flag=n
```

This would become:

```
switch type=string flag=nh
```

5.8.6 Update Fairshare Entities and Database

This step is only necessary if you were not using `egroup:euser` as your fairshare entity, but wish to do so now. Follow the procedures in section 5.10.29, “Updating Fairshare Definitions for `egroup:euser`”, on page 132.

5.8.7 Start the New PBS

1. Start PBS on the execution hosts. On each machine, type:

```
PBS_EXEC/etc/pbs_habitat
```

```
PBS_EXEC/sbin/pbs_mom -p
```

2. Start PBS on the server’s host:

```
<path to init.d>/init.d/pbs start
```

5.8.8 Set the Server’s `pbs_license_info` Attribute

Set the server’s `pbs_license_info` attribute to point to the license file or license server. See section 4.4.3.1, “Setting the Licensing Information in `pbs_license_info`”, on page 87.

5.9 Platform-Specific UNIX/Linux Overlay Upgrades

5.9.1 Upgrading on an Altix or a Complex Containing One or More Altixes

This section contains instructions for an overlay upgrade of an Altix on which PBS will manage the cpusets. If you want to configure PBS on the Altix to manage cpusets, run `pbs_mom.cpuset`. Jobs cannot be running on an Altix during the upgrade. Jobs on the Altix can be requeued, killed, or allowed to finish running. Multi-vnode jobs cannot survive an overlay upgrade. A multi-vnode job is killed and requeued when any of its MOMs is restarted.

If you want PBS to manage cpusets after the upgrade, you must have a vnode definitions file. The vnode definitions file is generated automatically for an Altix running ProPack 6 or 7 or Performance Suite 1.

5.9.1.1 Back Up Your Existing PBS Professional

Make a tarfile of the PBS_HOME and PBS_EXEC directories.

1. Make a backup directory:

```
mkdir /tmp/pbs_backup
```

2. Make a tarfile of PBS_HOME:

```
cd PBS_HOME/..  
tar -cvf /tmp/pbs_backup/PBS_HOME_backup.tar PBS_HOME
```

3. Make a tarfile of PBS_EXEC:

```
cd PBS_EXEC/..  
tar -cvf /tmp/pbs_backup/PBS_EXEC_backup.tar PBS_EXEC
```

4. Make a copy of your configuration file:

```
cp /etc/pbs.conf /tmp/pbs_backup/pbs.conf.backup
```

5. If they exist (PBS 8.0 and later), make a copy of the Version 2 configuration files on each machine:

```
mkdir /tmp/pbs_backup/mom_configs  
$PBS_EXEC/sbin/pbs_mom -s list | egrep -v '^PBS' | while read file  
do  
    $PBS_EXEC/sbin/pbs_mom -s show $file > /tmp/pbs_backup/mom_configs/  
    $file  
done
```

6. Save a list of the hosts:

```
pbsnodes -a > /tmp/pbs_backup/hostlist
```

7. Save the server's node configuration:

```
Qmgr: print nodes @default > /tmp/pbs_backup/newnodes
```

8. If you are upgrading from a pre-8.0 version, ensure that each Altix host has its values for resources_available.(mem|vmem|ncpus) unset:

```
Qmgr: unset node <hostname> resources_available.mem  
Qmgr: unset node <hostname> resources_available.ncpus  
Qmgr: unset node <hostname> resources_available.vmem
```

9. Make a copy of the scheduler's directory to modify:

```
cp PBS_HOME/sched_priv /tmp/pbs_backup/sched_priv.work
```

5.9.1.2 Stop New Jobs From Starting

You must stop any jobs from starting. Do this by stopping scheduling.

Stop scheduling:

```
Qmgr: set server scheduling=false
```

5.9.1.3 Stop Jobs Running on the Altix

Jobs cannot be running on an Altix during an upgrade from a version before 8.0 to 8.0 or later. You can do one of the following:

- requeue any jobs running on the Altix
- kill the jobs on the Altix
- drain the Altix by letting existing jobs on the Altix finish running
- requeue all Altix and non-Altix jobs in the complex

If you choose to requeue all Altix and non-Altix jobs, you can skip to section 5.9.1.4, “Shut Down Your Existing PBS”, on page 115. This section has instructions for requeueing all jobs in the complex. Note that if you choose to requeue jobs, those jobs that are marked non-runnable will be killed.

To requeue only jobs running on the Altix:

1. List the jobs on the Altix This will list some jobs more than once. You only need to requeue each job once:

```
pbsnodes <hostname> | grep Jobs
```

2. Requeue the jobs:

```
qrerun <job ID> <job ID> ...
```

To kill the jobs on the Altix:

1. List the jobs on the Altix This will list some jobs more than once. You only need to kill each job once:

```
pbsnodes <hostname> | grep Jobs
```

2. Use the `qdel` command to kill each job by job ID:

```
qdel <job ID> <job ID> ...
```

To drain the Altix, wait until any jobs running on the Altix have finished.

5.9.1.4 Shut Down Your Existing PBS

You can let any non-Altix jobs continue to run, or you can requeue or kill all jobs.

To let running jobs continue to run on non-Altix hosts:

Shut down PBS. The `qterm` command will use the `-t quick` option unless you specify otherwise.

(PBS versions prior to PBSPro_5.4.0)

```
qterm -m -s
```

(PBS versions PBSPro_5.4.0 and later)

```
qterm -m -s -f
```

If your server is not running in a failover environment, the “`-f`” option is not required.

To requeue or kill all jobs:

Shut down PBS. The `qterm -t immediate` command will requeue any jobs that can be requeued and kill those that cannot:

```
qterm -t immediate -m -s -f
```

If your server is not running in a failover environment, the “`-f`” option is not required.

5.9.1.5 Install the New Version of PBS

On any non-Altix platforms, follow the instructions in section 5.8.3, “Install the New Version of PBS”, on page 111.

On an Altix, follow the instructions in section 3.4.8.2, “Download and Install the New PBS”, on page 48.

5.9.1.6 Prepare the New Scheduler's Configuration File

1. Make a copy of the new `sched_config`, which is in `PBS_EXEC/etc/pbs_sched_config`.

```
cp PBS_EXEC/etc/pbs_sched_config PBS_EXEC/etc/pbs_sched_config.new
```
2. Update `PBS_EXEC/etc/pbs_sched_config.new` with any modifications that were made to the current `PBS_HOME/sched_priv/sched_config`.
3. If it exists, replace the `strict_fifo` option with `strict_ordering`. If you do not, a warning will be printed in the log when the scheduler starts.
4. If you are upgrading from a version prior to 7.1 and copied over your scheduler log filter setting, make sure the new configuration file has 1024 added to it. If the value is less than 1024, add 1024 to it:

Edit `PBS_EXEC/etc/pbs_sched_config.new`.

If your previous log filter line was:

```
log_filter:256
```

change it to:

```
log_filter:1280
```

If you do not, you will be inundated with logging messages.

5. Move `PBS_EXEC/etc/pbs_sched_config.new` to the correct name and location, i.e. `PBS_HOME/sched_priv/sched_config`.

```
mv PBS_EXEC/etc/pbs_sched_config.new PBS_HOME/sched_priv/sched_config
```

5.9.1.7 Update Fairshare Entities and Database

This step is only necessary if you were not using `egroup:euser` as your fairshare entity, but wish to do so now. Follow the procedures in section 5.10.29, "Updating Fairshare Definitions for `egroup:euser`", on page 132.

5.9.1.8 Modify the New Server's Resource File

If you are upgrading from a version prior to 7.1, add the "h" flag to those vnode-level resources listed in the server's `PBS_HOME/server_priv/resourcedef` file that have the "n" or "f" flag. For example, if you had:

```
switch type=string flag=n
```

This would become:

```
switch type=string flag=nh
```

See section 5.14.2.5, "Resource Accumulation Flags" on page 295 in the PBS Professional Administrator's Guide.

5.9.1.9 Edit the Altix Configuration File

When changing from ProPack 2 or 3 to Propack 6 or 7 or Performance Suite 1, remove any `cpuset_create_flags <flags>` initialization other than `CPUSET_CPU_EXCLUSIVE` from the Version 1 MOM config file. See the `pbs_mom(8B)` manual page.

5.9.1.10 Start the New PBS

1. On any non-Altix execution hosts, start PBS. On each of these machines, type:

```
PBS_EXEC/etc/pbs_habitat
```

```
PBS_EXEC/sbin/pbs_mom -p
```

2. On any Altixes, start PBS. For the location of the startup script, see section 6.1, "Starting and Stopping PBS: UNIX and Linux", on page 167.

```
<path to script>/pbs start
```

3. If the server's host is not an Altix, start the scheduler and server using the following commands:

```
<path to init.d>/init.d/pbs start
```

5.9.1.11 Read the Nodes Information Back Into `qmgr`

Read the nodes information back into `qmgr`:

```
qmgr < /tmp/pbs_backup/newnodes
```

5.9.1.12 Set the Server's `pbs_license_info` Attribute

Set the server's `pbs_license_info` attribute to point to the license file or license server. See section 4.4.3.1, "Setting the Licensing Information in `pbs_license_info`", on page 87.

5.9.1.13 Define Resources for the Altix

If you are upgrading an Altix from a pre-8.0 version of PBS Professional, you must change how resources are defined. The script “`au-nodeupdate.pl`” does this for you. It takes the list of hosts as an argument, and defines each resource on the host's natural vnode, then defines the resources indirectly on each of the hosts' vnodes. It does this for each resource defined on the host that is not `mem`, `vmem` or `ncpus`.

```
PBS_EXEC/etc/au_nodeupdate.pl /tmp/pbs_backup/hostlist
```

5.10 Migration Upgrade Under UNIX and Linux

Follow these instructions if you are upgrading a Cray or an Altix on which PBS will **not** be managing the cpusets, that is, you will run the standard, non-cpuset MOM, or if you are upgrading from a much older version of PBS Professional. See the PBS Professional Release Notes for recommendations on which versions require a migration upgrade.

`PBS_EXEC` is automatically created as a versioned directory. The installation program creates the directory for you. Each version of the PBS binaries is installed in its own versioned directory, so when you migrate from one version to another, you do not need to move the old version of PBS. You can let any jobs complete execution.

You will probably want to move jobs from the old PBS to the new. During a migration upgrade, jobs cannot be running. You can checkpoint, terminate and requeue all possible jobs and requeue non-checkpointable but rerunnable jobs. Your options with non-rerunnable jobs are to either let them finish or kill them.

The migration upgrade instructions here are generic. During the installation step, follow the directions for your operating system.

In the instructions below, file and directory pathnames are the PBS defaults. If you installed PBS in different locations, use your locations instead. `PBS_EXEC_OLD` refers to your existing, pre-upgrade location for `PBS_EXEC`.

The following commands must be run as “root”.

5.10.1 Back Everything Up

5.10.1.1 Back Up Server/scheduler Host

Back up the server and vnode configuration information. You will use it later in the migration process. Do the following on the server/scheduler host:

1. On the server host, create a backup directory called `/tmp/pbs_backup`

```
mkdir /tmp/pbs_backup
```

2. Print the server attributes to a backup file in the backup directory:

```
qmgr -c "print server" > /tmp/pbs_backup/server.backup
```

3. Make a copy of the server's configuration for the new PBS:

```
cp /tmp/pbs_backup/server.backup /tmp/pbs_backup/server.new
```

4. Print the vnode attributes and creation commands for the default server to a backup file in the backup directory. The default server is specified in `/etc/pbs.conf`.

```
qmgr -c "print node @default" > /tmp/pbs_backup/nodes.backup
```

5. Make a copy of the vnode attributes for the new PBS:

```
cp /tmp/pbs_backup/nodes.backup /tmp/pbs_backup/nodes.new
```

6. Make a copy of the server's `resourcedef` file for the new PBS:

```
cp PBS_HOME/server_priv/resourcedef /tmp/pbs_backup/resourcedef.new
```

7. Make a copy of the scheduler's `sched_priv` directory for the new PBS:

```
cp -rp PBS_HOME/sched_priv/ /tmp/pbs_backup/sched_priv.work
```

8. Copy the `pbs.conf` file to the backup directory:

```
cp /etc/pbs.conf /tmp/pbs_backup/pbs.conf.backup
```

5.10.1.2 Back Up Execution Hosts

On each execution host, back up the MOM configuration files:

1. Make a copy of the Version 1 configuration file:

```
cp PBS_HOME/mom_priv/config /tmp/pbs_backup/config.backup
```

2. If they exist (PBS 8.0 and later), make a copy of the Version 2 configuration files:

```
mkdir /tmp/pbs_backup/mom_configs
$PBS_EXEC_OLD/sbin/pbs_mom -s list | egrep -v '^PBS' | while read file
do
    $PBS_EXEC_OLD/sbin/pbs_mom -s show file > /tmp/pbs_backup/mom_configs/
    $file
done
```

Make a tarfile of the PBS_HOME and PBS_EXEC_OLD directories. This is a precaution.

1. Make a tarfile of PBS_HOME:

```
cd PBS_HOME/..
tar -cvf /tmp/pbs_backup/PBS_HOME_backup.tar PBS_HOME
```

2. Make a tarfile of PBS_EXEC_OLD:

```
cd PBS_EXEC_OLD/..
tar -cvf /tmp/pbs_backup/PBS_EXEC_OLD_backup.tar PBS_EXEC_OLD
```

3. Make a backup of the existing /etc/pbs.conf:

```
cp /etc/pbs.conf /tmp/pbs_backup/pbs.conf.backup
```

5.10.2 Replace Properties with Boolean Resources, Update Resources

You must replace properties with Boolean resources.

1. Manually edit the vnodes configuration file for the new PBS, `/tmp/pbs_backup/nodes.new`. Where you find a line of the form:

```
set node NAME properties=PROP
```

or

```
set node NAME properties+=PROP
```

Replace with the line:

```
set node NAME resources_available.PROP=True
```

For example, if the `qmgr` output contained the following lines:

```
set node node01 properties=red
set node node01 properties+=green
set node node02 properties=red
set node node02 properties+=blue
```

replace those lines with:

```
set node node01 resources_available.red=True
set node node01 resources_available.green=True
set node node02 resources_available.red=True
set node node02 resources_available.blue=True
```

2. Create Boolean resources to replace properties. For each property being replaced, create or append to the new `/tmp/pbs_backup/resourcedef.new` file a line of the form:

```
PROP type=boolean flag=h
```

In the previous step's example, you would add the following lines to `/tmp/pbs_backup/resourcedef.new`:

```
red type=boolean flag=h
green type=boolean flag=h
blue type=boolean flag=h
```

You only need to add each former property once.

3. Add the "h" flag to the "n" or "f" flag for vnode-level resources listed in the new server's `/tmp/pbs_backup/resourcedef.new` file. For example, if you had:

```
switch type=string flag=n
```

This would become:

```
switch type=string flag=nh
```

5.10.3 Remove Deprecated Terms From Server and Vnode Configurations

1. On the server host, manually edit the vnodes configuration file for the new PBS, `/tmp/pbs_backup/nodes.new`. Delete all occurrences of:

```
ntype=cluster
```

or

```
ntype=time-shared
```

Otherwise you will get a harmless error.

2. Manually edit the server's configuration file for the new PBS, `/tmp/pbs_backup/server.new`. Delete all lines of the form:

```
resources_default.neednodes=X
```

5.10.4 Preserve Reservation Information

After the upgrade, you will have to re-create any reservations that exist on your old server. You can print all of your reservation information to a file.

Print the reservation information to the file `/tmp/pbs_backup/reservations`.

```
pbs_rstat -f > /tmp/pbs_backup/reservations
```

5.10.5 Prevent Jobs From Being Enqueued or Started

You must deactivate the scheduler and queues. When the server's `scheduling` attribute is false, jobs are not started by the scheduler. When the queues' `enabled` attribute is false, jobs cannot be enqueued.

1. Prevent the scheduler from starting jobs:

```
qmgr -c "set server scheduling = false"
```

2. Print a list of all queues managed by the server. Save the list of queue names for the next step.

```
qstat -q
```

3. Disable queues to stop jobs from being enqueued. Do this for each queue in your list from the previous step.

```
qdisable <queue name>
```

5.10.6 Allow Running Jobs to Finish, or Requeue Them

You cannot perform a migration upgrade while jobs are running. Either let running jobs finish, or requeue them. (You can also delete them.)

To requeue any jobs running:

1. List the jobs. This will list some jobs more than once. You only need to requeue each job once:

```
pbsnodes <hostname> | grep Jobs
```

2. Requeue the jobs:

```
qrerun <job ID> <job ID> ...
```

To kill the jobs:

1. List the jobs. This will list some jobs more than once. You only need to kill each job once:

```
pbsnodes <hostname> | grep Jobs
```

2. Use the `qdel` command to kill each job by job ID:

```
qdel <job ID> <job ID> ...
```

To drain the host, wait until any running jobs have finished.

5.10.7 Shut Down PBS Professional

You can now shut down the server, scheduler, and MOM daemons. Use the `-t immediate` option to `qterm` so that all possible running jobs will be queued.

1. Shut down PBS:

```
qterm -t immediate -m -s -f
```

If your server is not running in a failover environment, the “-f” option is not required.

2. Verify that PBS daemons are not running in the background:

```
ps -ef | grep pbs
```

If you see the `pbs_server`, `pbs_sched`, or `pbs_mom` process running, you will need to manually terminate that process:

```
kill -9 <pid>
```

5.10.8 Back Up the Server’s Jobs Directory

You must back up any jobs that were queued when the server was shut down, in order to move them to the new version of PBS. Do the following on the server host:

1. Make a tarfile of the jobs directory, and save it in the backup directory:

```
cd PBS_HOME/server_priv  
tar -cvf /tmp/pbs_backup/pbs_jobs_save.tar jobs
```

5.10.9 Back Up PBS Directories and Configuration Files

Back up the `PBS_HOME` directory and PBS configuration files. You will use this later.

1. Rename the `PBS_HOME` directory:

```
mv PBS_HOME PBS_HOME.backup
```

5.10.10 Install the New Version of PBS

For Cray, follow the instructions in section 5.11.1, “Installing PBS To Upgrade on the Cray XT”, on page 134.

For other UNIX/Linux platforms:

1. On each host, go to the directory containing the package for the new version of PBS. Unzip and untar or otherwise unpack the new version of PBS. Install the new version of PBS.

```
cd <location of package>
tar -xvf <pbs package>
./INSTALL
```

2. If that machine will be the server and scheduler host, select option #1.
3. When asked for the license file location, give the location of the license file.
4. When asked whether to start PBS, DO NOT start PBS daemons. You will manually start the server later.
5. Create the subdirectories under PBS_HOME. On the new PBS server host and on each execution host:

```
{PBS_EXEC}/etc/pbs_habitat
```

5.10.11 Transfer MOM Configuration Information

On each execution host (for Cray, login node), follow these steps:

1. Restore the Version 1 MOM configuration file.

```
cp /tmp/pbs_backup/config.backup PBS_HOME/mom_priv/config
```

2. If they exist (PBS 8.0 and later), restore the Version 2 MOM configuration files. For each of these that you backed up:

```
cd /tmp/pbs_backup/mom_configs
for file in *
do
  $PBS_EXEC/sbin/pbs_mom -s insert $file $file
done
```

5.10.12 Update Fairshare Entities and Database

This step is only necessary if you were not using `egroup:euser` as your fairshare entity, but wish to do so now. Follow the procedures in section 5.10.29, “Updating Fairshare Definitions for `egroup:euser`”, on page 132.

5.10.13 Copy License and Resource Files

3. Copy the modified server resource file from the backup directory to the default directory:

```
cp /etc/pbs_backup/resourcedef.new PBS_HOME/server_priv/resourcedef
```

5.10.14 Start the New Server Without Defined Queues or Vnodes

When the new server starts up it will have default queue “workq” and the server host already defined. You want to start the new server with empty configurations so that you can import your old settings.

1. Start the new server with empty queue and vnode configurations:

```
PBS_EXEC/sbin/pbs_server -t create
```

A message will appear saying “Create mode and server database exists, do you wish to continue?”

Type “**y**” to continue.

Because of the new licensing scheme an additional message appears:

```
"One or more PBS license keys are invalid, jobs may not run"
```

This message is expected. Continue to the next step in these instructions.

5.10.15 Set Server’s License Location Attribute

Follow the steps in section 4.4.3.1, “Setting the Licensing Information in pbs_license_info”, on page 87

5.10.16 Replicate Queues and Server and Vnodes Configuration

1. Give the new server the old server's configuration, but modified for the new PBS:

```
PBS_EXEC/bin/qmgr < /tmp/pbs_backup/server.new
```

2. Replicate vnodes configuration, also modified for the new PBS:

```
PBS_EXEC/bin/qmgr < /tmp/pbs_backup/nodes.new
```

3. Verify the original configurations were read in properly:

```
PBS_EXEC/bin/qmgr -c "print server"
```

```
PBS_EXEC/bin/pbsnodes -a
```

5.10.17 Re-create Reservations

You must re-create each reservation that was on the old server, using the `pbs_rsub` command. Each reservation is created as a new reservation. You can use all of the information about the old reservation except for its start time. Be sure to give each reservation a start time in the future. Use the information stored in `/tmp/pbs_backup/reservations`.

5.10.18 Change Port Specifications in `pbs.conf` for Old PBS

You must edit the `pbs.conf` file of the old PBS so that all old services use ports that won't clash with those of the new PBS. The old `pbs.conf` file was saved as `/etc/pbs.conf.pre.M.N.P.S`, where M.N.P.S is the version of the new PBS.

You must change the port numbers for all PBS related daemons: server, MOM, scheduler, and data service.

For example, if the old version of PBS is 11.0.2.110766, and the new one is 11.1.0.111247, the name of the old PBS configuration file is `/etc/pbs.conf.pre.11.1.0.111247`.

You edit this file so that the entries look like those in the following table:

Table 5-1: Entries in Old PBS Configuration File

New Entry in <code>pbs.conf</code>	Description
<code>PBS_EXEC=/opt/pbs/11.0.2.110766</code>	This is the path for the old <code>PBS_EXEC</code>
<code>PBS_HOME=/var/spool/PBS.old</code>	This is the path for the old <code>PBS_HOME</code>

Table 5-1: Entries in Old PBS Configuration File

New Entry in <code>pbs.conf</code>	Description
<code>PBS_START_SERVER=1</code>	Unchanged
<code>PBS_START_MOM=1</code>	Unchanged
<code>PBS_START_SCHED=1</code>	Unchanged
<code>PBS_SERVER=x17-64-rhel6</code>	Unchanged
<code>PBS_BATCH_SERVICE_PORT= 13001</code>	This is the changed port number for the old server
<code>PBS_DATA_SERVICE_PORT=13007</code>	This is the changed port number for the old data service
<code>PBS_MOM_SERVICE_PORT= 13002</code>	This is the changed port number for the old MOM
<code>PBS_SCHEDULER_SERVICE_PORT= 13003</code>	This is the changed port number for the old scheduler

5.10.19 Set Path to Old `pbs.conf`

Set the `PBS_CONF_PATH` environment variable to point to the old PBS configuration file:

```
export PBS_CONF_FILE=/etc/pbs.conf.pre.11.1.0.111247
```

5.10.20 Start the Old Server

You must start the old server in order to move jobs to the new server. The old server must be started on **alternate** ports. These are specified in the configuration file for the old PBS.

Start the old server daemon:

```
PBS_EXEC_OLD/sbin/pbs_server
```

5.10.21 Unset Path to Old `pbs.conf`

Unset the `PBS_CONF_PATH` environment variable:

```
unset PBS_CONF_FILE
```

5.10.22 Move Existing Jobs to the New Server

You must move existing jobs from the old server to the new server. To do this, you run the move commands from the old server, and give the new server's port number, 15001, in the destination. See the `qmove(1B)` man page. When moving jobs from reservation queues, be sure to move them into the equivalent new reservation queues.

This is for the special case where your old PBS version is older than 5.4.0 and the old server's host also runs a MOM:

1. Delete the vnode on the server's host:

```
PBS_EXEC_OLD/bin/qmgr -c "d n <old server host>" <old server host>:13001
```

If you see the message, "Cannot delete busy object":

- a. Get a list of jobs running on that vnode:

```
PBS_EXEC_OLD/bin/qstat @<old server host>:13001
```

- b. Either requeue or kill the jobs on the server's host:

```
PBS_EXEC_OLD/bin/qrerun -W force <job id>
```

For PBS versions 5.4.0 and later, if your old server's host also ran a MOM, you will need to delete that vnode from the old server.

Delete the vnode on the old server's host:

```
PBS_EXEC_OLD/bin/qmgr -c "d n <old server host>" <old server host>:13001
```

Move jobs from the old server to the new one:

1. Print the list of jobs on the old server:

```
PBS_EXEC_OLD/bin/qstat @<old server host>:13001
```

2. Move each job from each queue:

```
PBS_EXEC_OLD/bin/qmove <new queue name>@<new server host>:15001 <job id>@<old server host>:13001
```

You can use `qselect` to select all the jobs within a queue instead of moving each job individually.

3. Move all jobs within a queue:

```
PBS_EXEC_OLD/bin/qmove <queue name>@<new server host>:15001 `PBS_EXEC_OLD/bin/qselect -q <queue name>@<old server host>:13001`
```

If you see the error message "Too many arguments...", there are too many jobs to fit in the shell's command line buffer. You can continue moving jobs one at a time until there are few enough.

5.10.23 Shut Down Old Server

Shut down the old server daemon:

```
PBS_EXEC_OLD/bin/qterm -t quick <old server host>:13001
```

5.10.24 Update New `sched_config`

Update the new scheduler's configuration file, in `PBS_HOME/sched_priv/sched_config`, with any modifications that were made to the old `PBS_HOME.old/sched_priv/sched_config`.

1. If you copied over your old scheduler log filter value, make sure that it has had 1024 added to it. If the value is less than 1024, add 1024 to it. For example, if the old log filter line is:
`log_filter:256`
change it to:
`log_filter:1280`
2. If it exists, replace the `strict_fifo` option with `strict_ordering`. If you do not, a warning will be printed in the log when the scheduler starts.

5.10.25 Start New Scheduler

Start the scheduler daemon, on the server's host:

```
PBS_EXEC/sbin/pbs_sched
```

5.10.26 Start New MOMs

Run the start script on each execution host, so that the subdirectories for `PBS_HOME` are created. On each execution host:

```
<path to init.d>/init.d/pbs start
```

5.10.27 Optionally Start MOM on New Server's Host

If your old configuration had a MOM running on the server's host, and you wish to replicate the configuration, you can start a MOM on that machine.

Start the MOM daemon on the new server's host:

```
PBS_EXEC/sbin/pbs_mom
```

5.10.28 Enable Scheduling in New Server

You must set the new server's scheduling attribute to true so that the scheduler will start jobs.

Enable scheduling for the new server:

```
PBS_EXEC/bin/qmgr -c "s s scheduling=1"
```

5.10.29 Updating Fairshare Definitions for egroup:euser

This step is only necessary if you were not using egroup:euser as your fairshare entity, but wish to do so now.

On the server/scheduler host:

1. Update /tmp/pbs_backup_sched_priv.work/resource_group, with the new egroup:euser entities.

If you previously had an euser named Bob, the new entry for Bob in the resource_group file will look like this:

```
pbsgroup1:Bob          101    Phys          20
```

2. Restore the modified contents of the scheduler's directory:

```
cp -r /tmp/pbs_backup/sched_priv.work PBS_HOME/sched_priv
```

If you wish to keep your old fairshare resource usage data, use the pbsfs command to update the data. See the pbsfs (8B) manual page.

When you display the usage data, both the old and the new entities will appear. Entities defined before the upgrade will appear in the unknown group. Entities defined in the previous step will be in the correct department but have no usage.

1. Display the usage data:

```
% pbsfs -p
```

If the old entity was Bob:

```
Bob : group: 1 cgrp: 5 shares: 10 Usage: 1005 Perc: 50%
```

You will see usage of 1005 for Bob, as shown above. The newly defined entity is `pbsgroup1:Bob`, and it will appear like this:

```
pbsgroup1:Bob 5 4 root 10
```

2. Assign Bob's usage to `pbsgroup1:Bob`:

```
% pbsfs -s pbsgroup1:Bob 1005
```

3. Remove the old entities:

```
% pbsfs -e
```

If your previous entities were `euser`, and now you have `egroup:euser`, and you have a user who is in more than one group, you will need to make an entry in `/tmp/pbs_backup_sched_priv.work/resource_group` for Bob in each of his groups, then assign a portion of the Bob's usage to each group.

For example, if you wish to have Bob be in two groups, `pbsgroup1` and `pbsgroup2`, you may want Bob's usage to be split between the groups. Bob will show up as `pbsgroup1:Bob` and `pbsgroup2:Bob`.

1. Update `/tmp/pbs_backup_sched_priv.work/resource_group`, putting Bob into `pbsgroup1` and `pbsgroup2`.

The entries for Bob in the `resource_group` file will look like this:

```
pbsgroup1:Bob          101    Phys          20
pbsgroup2:Bob          102    Math          20
```

2. Assign half of Bob's usage to each of `pbsgroup1:Bob` and `pbsgroup2:Bob`:

```
% pbsfs -s pbsgroup1:Bob 502
```

```
% pbsfs -s pbsgroup2:Bob 503
```

5.11 Upgrade Installation Instructions

5.11.1 Installing PBS To Upgrade on the Cray XT

For the Cray, a migration upgrade is required.

These are the instructions for the software installation process only. In order to perform an upgrade of PBS, you must follow the complete set of migration upgrade instructions. See section 5.10, “Migration Upgrade Under UNIX and Linux”, on page 119.

The installation/upgrading procedures here are new. Please read this section carefully.

As of version 11.0, PBS Professional uses licenses managed by a new Altair License Server. The license server **must** be installed and configured before PBS is installed. Begin with section 3.1, “Requirements for Installing PBS Professional”, on page 25.

The PBS Professional package supports CNL.

The PBS Professional Scheduler and Server can be run on an external host; they do not have to run on your Cray system. In the following instructions, the examples show the PBS Server and Scheduler running on the `sdb` node. One instance of the `pbs_mom` daemon is run on each login node of the Cray system, allowing PBS Professional to select the login node to which a job is assigned.

The PBS Professional installation script, `INSTALL`, will detect that it is running on a Cray XT and perform special actions to manage `PBS_EXEC`, `/etc/pbs.conf`, and `PBS_HOME`.

5.11.1.1 Not Supported

Kerberos is not supported on the Cray XT.

5.11.1.2 Required Libraries

The XML library, `libexpat.a`, must be available on the login nodes. This is fairly standard in Linux distributions.

5.11.1.3 Use Correct Version Number

You must know the version number of the PBS Professional package you are installing. The following examples use version number 10.0.0.08470. The version number of the package you are installing will be different. Please be sure to use the version number of your package.

5.11.1.4 Hardware and Software Requirements

Confirm the following before beginning the upgrade:

- You must have root permission.
- All configuration files (such as `/etc/nsswitch.conf`) must have been updated appropriately for your site. The `/etc/nsswitch.conf` file must contain the following lines:

```
hosts: files dns
networks: files dns
```

If the `/etc/nsswitch.conf` file is not updated, the upgrade will fail and you will receive the following error message:

```
pbsd_main: unable to get my hostname
```

- Verify that `nfslock` and `rstatd` are on:

```
# ssh root@boot
boot:~# xtopview
default/:// # chkconfig
rstatd on
nfslock on
default/:// # exit
```

- The PBS Server and Scheduler must be installed on a node with network access to the Altair License Server. This will allow the PBS Server to use the PBS Professional licenses. This may be the `sdb` node on the XT or it may be a node separate from the XT system.
- On the login nodes and on the node on which the PBS Server and Scheduler will run (typically the `sdb` node), make sure that the `/var` directory is owned by root.
- Determine how you will start up and shut down PBS Professional. You must choose whether to configure resiliency communication agent (RCA) services or use `init.d` (`chkconfig`) to start up PBS Professional.
- Determine whether you need to set up a License Server or request a new PBS Professional license key from Altair Engineering. For more information on generating and installing license keys, see the Altair License Management System Installation and Operations Guide, available on the Altair website at www.altair.com.
- It is recommended that you back up your current configuration. If the upgrade process should fail unexpectedly, you can restore your existing server database and configuration files. Log in to the boot node as root, and then log in to the `sdb` node. Archive the existing

directories.

```
smw:~# ssh root@boot
```

```
boot001:~# ssh root@sdb
```

```
sdb:~# tar -cvzf /var/PBS-host_name.tar.gz /var/spool/PBS
```

- Compute nodes must be in the “batch role” before they can be used by jobs under the control of PBS Professional. To set all of the compute nodes in the batch role:

On the boot node execute:

```
xtprocadm -k m batch
```

Restart ALPS:

```
/etc/init.d/alps restart
```

- Verify that the machine is running CLE 2.2 or higher.

5.11.1.5 Installing the New Version of PBS

To install the new version of PBS Professional, follow the steps listed below:

1. Obtain the current distribution of PBS Professional. This is a gzipped tar file, for example:

```
PBSPro_10.0.0-linux24_x86_64_crayxt.tar.gz
```

2. Create a temporary install directory on the boot node:

```
smw:~# ssh root@boot mkdir /rr/current/software/pbs-install
```

3. Copy the distribution file to that directory:

```
smw:~# scp PBSPro_10.0.0-linux24_x86_64_crayxt.tar.gz root@boot:/rr/current/software/pbs-install
```

4. Log in as root to the boot node.

5. Change to the temporary install directory.

6. Unzip and untar the distribution:

```
smw:~# ssh root@boot boot001:~  
# cd /rr/current/software/pbs-install  
# gunzip PBSPro_10.0.0-linux24_x86_64_crayxt.tar.gz  
# tar xf PBSPro_10.0.0-linux24_x86_64_crayxt.tar
```

7. Change into the directory created when the package was untarred:

```
cd PBSPro_10.0.0
```

8. Run the `INSTALL` script in this directory:

```
# ./INSTALL
```

9. The install script checks to see whether you are using a supported version of ALPS. The script notifies you whether or not your version of ALPS is supported. If ALPS is not a supported version, the script prints a warning and asks whether to continue. You may

continue the installation, or you can halt it at this point. If you choose to continue, the behavior of PBS is not known.

10. Accept the offered default location for the installed binaries. The following is an example of how the installer offers a default location:

```
PBS_EXEC=PBSPro_11.0.0.08470
```

11. Accept the default location for the PBS working directories, for example:

```
PBS_HOME=/var/spool/PBS
```

12. When asked for the type of PBS Installation, choose 1:

1. Server, execution and commands.

13. When asked for the “License File Location(s)”, follow the instructions in section 3.3.4, “Setting PBS License Server Location”, on page 28.

14. You will be asked “Would you like to complete the install and switch to the new version of PBS?” To defer switching to the new version of PBS, answer “no”, and complete the upgrade later running the `pbs_postinstall` script as directed, in step 2 on page 140.

15. Verify that the reference copy of the new PBS Professional configuration file `/etc/pbs.conf.10.0.0.08470` exists in the default view of the shared root file system. The file should look like the following, except that “sdb” will be replaced with the PBS Server node you specified:

```
PBS_SERVER=sdb
PBS_EXEC=/opt/pbs/11.0.0.08470
PBS_HOME=/var/spool/PBS
PBS_START_SERVER=0
PBS_START_SCHED=0
PBS_START_MOM=0
```

This new configuration file will become the real `/etc/pbs.conf` during the post-install step later.

16. Remove the temporary install directory:

```
cd $HOME
rm -fr /rr/current/software/pbs-install
```

The new version of the PBS Professional software is now installed, and you are ready to proceed with configuration.

5.11.1.6 Configuring PBS Professional for the Cray XT

The new PBS Professional software must be installed before you configure it. The next steps guide you through the configuration process for the new version of PBS Professional.

5.11.1.7 Caveats

PATH must be included in the `pbs_environment` file. The PATH value is passed on to batch jobs. To maintain security, it is important that PATH be restricted to known, safe directories. Do not include "." in PATH.

5.11.1.8 Prerequisites to Configuration

If your site has disabled the use of remote operation functions ("r" commands) and output cannot be returned for jobs running on compute nodes, enable the use of the `cp` command by adding `$usecp` to the `$PBS_HOME/mom_priv/config` file on each login node.

5.11.1.9 Steps in Configuring PBS for the Cray XT

Review the following steps to ensure that your configuration has been retained:

1. As root, log in to the boot node

```
# ssh root@ boot
```

2. To complete the upgrade, use the `xtopview` command and run the post-install script. This step is only done if you answered “no” to the “Would you like to complete the install and switch to the new version of PBS?” when running the `INSTALL` script earlier.

```
# xtopview -m "postinstall" -e /opt/pbs/10.0.0.08470/etc/pbs_postinstall
```

Running `pbs_postinstall` updates `/etc/pbs.conf` with the new information and establishes `/opt/pbs/default` as a symbolic link to the new `PBS_EXEC`, which is, in this example, `/opt/pbs/10.0.0.08470`.

3. Tailor the install for the login nodes, enabling `pbs_mom`:

Use the `xtopview` command:

```
xtopview -m "pbs.conf for login" -c login
```

Run the `xtspec` command to specialize the change to `/etc/pbs.conf` for the login nodes:

```
class/login/: # xtspec /etc/pbs.conf
```

Edit `/etc/pbs.conf` ensuring the following lines are present:

```
PBS_START_MOM=1
```

Make sure that the Server and Scheduler do not start:

```
PBS_START_SERVER=0
```

```
PBS_START_SCHED=0
```

Ensure that the `PBS_SERVER` entry names the node on which you plan to run the Server and Scheduler, and that `PBS_EXEC` is set to `/opt/pbs/default`.

Ensure that `/opt/pbs/default` is a symbolic link to the location you specified for the binaries.

4. Exit from `xtopview`.
5. Determine the node ID, or NID, of the node on which you plan to run the PBS Server and Scheduler. This is usually the `sdb` node. Typically the NID of the `sdb` node is 3.

```
grep sdb /etc/hosts
```

6. Tailor the upgrade for the `sdb` node, enabling `pbs_server` and `pbs_sched`:

Use the `xtopview` command to run the `pbs_postinstall` script:

```
# xtopview -m "pbs.conf for Server" -n <NID>
```

Run the `xtspec` command to specialize the change to `/etc/pbs.conf` for the `sdb` node:

```
class/login/: # xtspec /etc/pbs.conf
```

Edit the `/etc/pbs.conf` file or verify that it contains the following lines:

```
PBS_START_SERVER=1
```

```
PBS_START_SCHEDULER=1
```

Verify that the MOM start line in `/etc/pbs.conf` is:

```
PBS_START_MOM=0
```

Ensure that the `PBS_SERVER` entry in `pbs.conf` names the node on which you plan to run the Server and Scheduler and that `PBS_EXEC` is set to `/opt/pbs/default`.

7. Exit from `xtopview`.

8. Log in as root to the node on which the PBS Server `pbs_server` will be running. This is shown here as the `sdb` node.

```
# ssh root@sdb
```

9. At this point, you **must** start the PBS Server with the `init.d/pbs` script. Do not start the server by any other means.

```
sdb # /etc/init.d/pbs start
```

This will create or update the PBS working directories, `PBS_HOME`, in `/var/spool`.

10. (Only necessary if you need to create new nodes) Using the `qmgr` command, create the PBS execution node or nodes, which are the login nodes, and set the special XT resource types for those nodes.

When creating a vnode to represent a login node, use the short name returned by the `hostname` command on the login node. For example, if `hostname` returns `HostA`, do the following:

```
Qmgr: create node HostA
```

If you create a vnode with a different name from the short name returned by `hostname`, the following happens:

- MOM creates a vnode whose name is the short name returned by `hostname`
- The vnode you created is not recognized by MOM, and is marked *stale*

Example of creating a vnode to represent a login node:

```
sdb: # module load pbs
```

```
sdb: # qmgr
```

```
Qmgr: create node login1
```

11. The default for `kill_delay` is 10 seconds. If the delay is less than that, it is recommended that you set it to 10 seconds.

```
sdb: # qmgr -c "set queue workq kill_delay=10"
```

12. Make sure that `flatuid` is set to true in the Server configuration:

```
sdb: # qmgr -c "list server"
```

If `flatuid` is not set to true, set it to true:

```
sdb: # qmgr -c "set server flatuid=true"
```

13. To ensure that scheduling of jobs is initiated when PBS is restarted, set the Server's scheduling attribute to true:

```
sdb: # qmgr -c "set server scheduling=true"
```

14. Shut down the PBS Services using the following command:

```
/etc/init.d/pbs stop
```

5.11.1.10 Configuration on Login Nodes

On each login node, log in and perform the following steps:

1. Start the PBS MOM daemon with the `init.d/pbs` script. Do not use any other means to start `pbs_mom`. This will create or update the PBS working directories, `PBS_HOME`, in `/var/spool`.

```
Login:~ # /etc/init.d/pbs start
```

2. Stop the MOM daemon by running the `init.d/pbs` script:

```
Login:~ # /etc/init.d/pbs stop
```

3. Ensure that the following directives are in the `pbs_mom` configuration file, `PBS_HOME/mom_priv/config`:

For CNL:

```
$clienthost sdb
$alps_client /usr/bin/apbasil
```

The sample configuration is now complete.

5.11.1.11 Configuring PBS Professional Startup and Shutdown

Do **NOT** start the new PBS Server yet.

Later, when it is time to start PBS, you can use `init.d` and the `chkconfig` command, or RCA services to start PBS Professional. To use `init.d` and the `chkconfig` command to start PBS Professional, execute Startup Procedure 1. To use RCA services to start PBS Professional execute Startup Procedure 2. These procedures are described in the following sections.

5.11.1.12 Startup Procedure 1: Using `init.d` (`chkconfig`) to Start PBS Professional

The PBS script is located in `/etc/init.d/pbs`. To use the `init.d` script and the `chkconfig` command to start up PBS Professional, execute the following steps:

1. Find the node ID (NID) of the node on which the PBS Server and Scheduler will run. If this is the `sdb` node, do this:

```
grep sdb /etc/hosts
```

2. Enable PBS startup for the Server and Scheduler on the `sdb` node. The NID of the `sdb` node is usually 3. This example uses an NID of 3.

```
boot001: # xtopview -n 3  
node/3:/ # chkconfig pbs on  
node/3:/ # exit
```

3. Enable PBS startup for the MOM on each login node:

```
# xtopview -c login  
class/login/: # chkconfig pbs on  
class/login/: # exit
```

5.11.1.13 Startup Procedure 2: Using RCA to Start PBS Professional

To configure RCA services to start PBS Professional, execute the following procedure:

1. The set of required files and directories that exist in the directory specified by `PBS_HOME` must be created or updated. `PBS_HOME` is normally `/var/spool/PBS`, but it may have been changed during the installation.

Run `{PBS_EXEC}/etc/pbs_habitat` for the Server/Scheduler node, which is usually the `sdb` node.

Run `{PBS_EXEC}/etc/pbs_habitat` for the login nodes.

2. Configure `pbs_mom` to start on each login node. Use the `xtservconfig` command:

```
sdb:~# /opt/xt-boot/default/bin/snos64/xtservconfig -c login add PBS-MOM
```

3. Configure `pbs_server` and `pbs_sched` to start on the appropriate node. For example, the following will run them on the `sdb` node (NID 3). See the `xtservconfig(8)` man page for more information.

```
sdb:~# /opt/xt-boot/default/bin/snos64/xtservconfig -n 3 add PBS-SERV
```

```
sdb:~# /opt/xt-boot/default/bin/snos64/xtservconfig -n 3 add PBS-SCHED
```

4. The scheduling process starts automatically when PBS Professional is started. To prevent this and have the scheduling start when the login nodes are booted, include the `-a false` option to the `SVC_START_CMD` line in the `/opt/xt-boot/default/etc/pbs-server` startup script. For example:

```
SVC_START_CMD="/opt/pbs/default/sbin/pbs_server -a false"
```

5.11.1.14 Transfer MOM Configuration Information

Note that it is not recommended to use vnode definition files on a Cray. Please see section 11.5.11.2, "Vnode Definition Files Not Recommended", on page 687 of the PBS Professional Administrator's Guide.

1. Restore the Version 1 MOM configuration file:

```
cp /tmp/pbs_backup/config.backup PBS_HOME/mom_priv/config
```

2. If they exist (PBS 8.0 and later), restore the Version 2 MOM configuration files. For each of these that you backed up:

```
cd /tmp/pbs_backup/mom_configs  
for file in *  
do  
  $PBS_EXEC/sbin/pbs_mom -s insert $file $file  
done
```

5.11.1.15 Check qstat Wrappers and Scripts

You must check any `qstat` wrappers or scripts. Either replace any `mpp*` resources, or modify the wrapper or script to use values in the job's `Submit_arguments` attribute. This attribute contains the original job submission line.

5.11.1.16 Replace Old `mpp*` Server and Queue Settings

As of PBS version 11.0, `mpp*` syntax is not supported. It is recommended that you replace `mpp*` settings for the server and queue. See section 11.5.7, "Recommended Manual Configuration", on page 675 of the PBS Professional Administrator's Guide for detailed instructions.

5.12 Upgrading Under Windows

You must use a migration upgrade under Microsoft Windows.

When you do a migration upgrade under Windows, you can install the new version of PBS in the same place or in a new location. You can install PBS in the default location or a non-default location.

You will probably want to move jobs from the old system to the new. During a migration upgrade, jobs cannot be running. You can requeue rerunnable jobs. You can let non-rerunnable jobs finish, or you can kill them.

If you are migrating from a version before 9.0 to 9.0 or later, all the following steps must be done by a domain administrator. If the migration is taking place in a domain environment, this Administrator account should be a member of the "Domain Admins" group. If the migration is taking place in a standalone environment, this Administrator account should be a member of the local "Administrators" group.

If you want the PBS service account not to be a domain administrator for the new version of PBS, follow the steps specified in section 5.12.28, "Optionally Change PBS Service Account to Non-domain Administrator Account", on page 166 before starting PBS.

In the instructions below, file and directory pathnames are the PBS defaults. If you installed PBS in different locations, use your locations instead. Where you see %WINDIR%, it will be automatically replaced by the correct directory. For Windows XP, that is \WINDOWS.

The default server is specified on 32-bit Windows systems in \Program Files\PBS Pro\pbs.conf and on 64-bit Windows systems in \Program Files (x86)\PBS Pro\pbs.conf.

Note that in version 8.0 and later, job scripts under Windows are executed differently. Any .bat files that are to be executed within a PBS job script will have to be prefixed with "call" as in:

```
---[job_b.bat]-----
    @echo off
    call E:\step1.bat
    call E:\step2.bat
-----
```

Without the "call", only the first .bat file gets executed, and it doesn't return control to the calling interpreter.

on 32-bit Windows systems, PBS is installed in \Program Files\PBS Pro\, and on 64-bit Windows systems, PBS is installed in \Program Files (x86)\PBS Pro\. In the following instructions, we use INSTALLATION_BASE to indicate \Program Files\ or \Program Files (x86)\ . When you follow the instructions, replace INSTALLATION_BASE with the appropriate directory.

5.12.1 Prevent Jobs From Being Enqueued or Started

You must deactivate the scheduler and queues. When the server's `scheduling` attribute is `false`, jobs are not started by the scheduler. When the queues' `enabled` attribute is `false`, jobs cannot be enqueued.

1. Prevent the scheduler from starting jobs:

```
qmgr -c "set server scheduling = false"
```

2. Print a list of all queues managed by the server. Save the list of queue names. You will need it in the next step and when moving jobs.

```
qstat -q
```

3. Disable queues to stop jobs from being enqueued. Do this for each queue in your list from the previous step.

```
qdisable <queue name>
```

5.12.2 Back Everything Up

Back up the server and vnode configuration information. You will use it later in the migration process.

1. Make a backup directory:

```
mkdir "%WINDIR%\TEMP\PBS Pro Backup"
```

2. Print the server attributes to a backup file in the backup directory:

```
qmgr -c "print server" > "%WINDIR%\TEMP\PBS Pro Backup\server.backup"
```

3. Make a copy of the server's configuration for the new PBS:

```
copy "%WINDIR%\TEMP\PBS Pro Backup\server.backup" "%WINDIR%\TEMP\PBS Pro Backup\server.new"
```

4. Print the default server's vnode attributes to a backup file in the backup directory.

```
qmgr -c "print node @default" > "%WINDIR%\TEMP\PBS Pro Backup\nodes.backup"
```

5. Make a copy of the vnode attributes for the new PBS:

```
copy "%WINDIR%\TEMP\PBS Pro Backup\nodes.backup" "%WINDIR%\TEMP\PBS Pro Backup\nodes.new"
```

6. Make a backup of the existing `INSTALLATION_BASE\PBS Pro\pbs.conf`. This command is all one line:

```
copy "INSTALLATION_BASE\PBS Pro\pbs.conf" "%WINDIR%\TEMP\PBS Pro Backup\pbs.conf.backup"
```

7. Make a copy of `pbs.conf` for the new PBS. This command is all one line:

```
copy "%WINDIR%\TEMP\PBS Pro Backup\pbs.conf.backup" "%WINDIR%\TEMP\PBS Pro Backup\pbs.conf.new"
```

8. Make a backup of the server's existing `resourcedef` file. This command is all one line:

```
copy "INSTALLATION_BASE\PBS Pro\home\server_priv\resourcedef" "%WINDIR%\TEMP\PBS Pro Backup\resourcedef.backup"
```

9. Make a copy of the server's `resourcedef` for the new PBS. This command is all one line:

```
copy "%WINDIR%\TEMP\PBS Pro Backup\resourcedef.backup" "%WINDIR%\TEMP\PBS Pro Backup\resourcedef.new"
```

10. Make a backup of the existing `INSTALLATION_BASE\PBS Pro\home\sched_priv`. This command is all one line:

```
copy "INSTALLATION_BASE\PBS Pro\home\sched_priv" "%WINDIR%\TEMP\PBS Pro Backup\sched_priv.backup"
```

11. Make a copy of the scheduler's directory for the new PBS. This command is all one line:

```
copy "%WINDIR%\TEMP\PBS Pro Backup\sched_priv.backup" "%WINDIR%\TEMP\PBS Pro Backup\sched_priv.work"
```

5.12.3 Preserve Reservation Information

After the upgrade, you will have to re-create any reservations that exist on your old server. You can print all of your reservation information to a file.

1. Print the reservation information to a file:

```
pbs_rstat -f > "%WINDIR%\TEMP\PBS Pro Backup\reservations"
```

5.12.4 Shut Down PBS Professional

You can now shut down the server, scheduler, and MOM daemons. Use the `-t immediate` option to the `qterm` command so that all possible running jobs will be queued.

1. Shut down PBS. If your server is not running in a failover environment, the `-f` option is not required.

```
qterm -t immediate -m -s -f
```

2. Stop the `pbs_rshd` daemon:

```
net stop pbs_rshd
```

5.12.5 Copy the Old Version of PBS to a Temporary Location

You will run the old PBS server from this temporary location in order to move jobs. You must do a copy rather than a move, because the installation software depends on the old version of PBS being available for it to remove.

On the Server vnode, copy the existing `PBS_HOME` and `PBS_EXEC` hierarchies to a temporary location. This command is all one line.

5.12.5.1 On Windows XP:

xcopy /o /e "\Program Files\PBS Pro" "%WINDIR%\TEMP\PBS Pro Backup"

Specify "D" for directory when prompted.

If you get an "access denied" error message while it is moving a file:

- a. Bring up Start menu->Programs->Accessories->Windows Explorer
- b. Right-click to select this file and bring up a pop-up menu.
- c. Choose "Properties", then "Security" tab, then "Advanced", then "Owners" tab.
- d. Reset the ownership of the file to "Administrators". "Administrators" must have permission to read the file.
- e. Rerun xcopy.

5.12.5.2 On Vista:

robocopy "\Program Files\PBS Pro" "%WINDIR%\TEMP\PBS Pro Backup" /S

If you get an "access denied" error message while it is moving a file:

- a. Bring up *Start menu->Programs->Accessories->Windows Explorer*
- b. Right-click on the folder/directory containing the problem files.
- c. Select *Properties|Security|Advanced*.
- d. Enable "*Allow inheritable permissions from parent to propagate to this object*" by checking it.
- e. Enable "*Reset permissions on all child objects and enable propagation of inheritable permissions*" by checking it.
- f. Click "*Apply*". A security box will advise you about propagation of inheritable permissions to child objects.
- g. Continue with "Yes".

Windows will then proceed to set or reset permissions on all child objects/folders/files. This may take a while, depending on the number of files in the directory tree, but the file progress is displayed with each file name. The files can now be read, copied or deleted without error.

Rerun robocopy.

5.12.6 Replace Properties with Boolean Resources and Update Resources

You must replace properties with boolean resources, and add a new flag for vnode-level resources.

1. Manually edit the vnodes configuration file for the new PBS, `%WINDIR%\TEMP\PBS Pro Backup\nodes.new`. You will find a line of the form:

```
set node NAME properties=PROP
```

or

```
set node NAME properties+=PROP
```

Replace with the line as a single line:

```
set node NAME resources_available.PROP=True
```

For example, if the `qmgr` output contained the following lines:

```
set node node01 properties=red
set node node01 properties+=green
set node node02 properties=red
set node node02 properties+=blue
```

replace those lines with these. Each is one line:

```
set node node01 resources_available.red=True
set node node01 resources_available.green=True
set node node02 resources_available.red=True
set node node02 resources_available.blue=True
```

2. Create Boolean resources to replace properties. For each property being replaced, create or append to the new `%WINDIR%\TEMP\PBS Pro Backup\resourcedef.new` file a line of the form:

```
PROP type=boolean flag=h
```

In the previous step's example, you would add the following lines to

`%WINDIR%\TEMP\PBS Pro Backup\resourcedef.new`:

```
red type=boolean flag=h
green type=boolean flag=h
blue type=boolean flag=h
```

You only need to add each former property once.

3. Add the "h" flag to the "n" or "f" flag for vnode-level resources listed in the new server's

`%WINDIR%\TEMP\PBS Pro Backup\resourcedef.new` file.

5.12.7 Remove Deprecated Terms from Server and Vnode Configurations

1. Manually edit the vnodes configuration file for the new PBS, `%WINDIR%\TEMP\PBS Pro Backup\nodes.new`. Delete all occurrences of:

`ntype=cluster`

or

`ntype=time-shared`

Otherwise you will get a harmless error.

2. Manually edit the server's configuration file for the new PBS, `%WINDIR%\TEMP\PBS Pro Backup\server.new`. Delete all lines of the form:

`resources_default.neednodes=X`

5.12.8 Install the New Version of PBS on Execution Hosts

You can install PBS from a CD or by downloading it.

If you are installing PBS from the CD, put it in the CD-ROM drive, browse to your CD-ROM drive, and click on the PBS Professional icon.

If you are installing PBS from a download, save it to your hard drive and run the self-extracting `pbspro.exe` package, either in the same directory, or by giving the path to it.

You must use the same password on all hosts. Do the following on each execution host, **except** the server's host.

Uninstall the old version of PBS:

1. Change your current directory so that it is not within `C:INSTALLATION_BASE\PBS Pro`, and make sure there is no access occurring on any file in that hierarchy. Otherwise you will have to remove the hierarchy by hand.
2. Run the installation program by either clicking on the icon, or typing:

PBSpro_<version>-windows.exe

Under XP, SP2, you may see a warning saying, "The publisher could not be verified. Are you sure you want to run this software?" Ignore this message and click the "Run" button.

The installation package will ask whether you want to uninstall the previous version.

Answer *yes*.

If you see a popup window saying the hierarchy could not be removed, remove the hierarchy manually by going to a command window and typing the following. Do not use the `del` command.

rmdir /S /Q "C:INSTALLATION_BASE\PBS Pro"

3. Reboot the execution host.

Install the new version of PBS:

1. Run the installation program by either clicking on its icon, or typing:

PBSpro_<version>-windows.exe

Under XP, SP2, you may see a warning saying, "The publisher could not be verified. Are you sure you want to run this software?" Ignore this message and click the "*Run*" button.

2. Check the installation location. You can change it, as long as the new location meets the requirements given in section 2.5, "Recommended PBS Configurations for Windows", on page 13.
3. Choose the "*Execution*" option.
4. Enter a non-empty password twice for the Administrator account. If you see "error 2245", check the password's length, complexity and history requirements. Look in the Active Directory guide or the Windows help page.
5. Give the server's hostname in the "Editing PBS.CONF file" window.
6. In the "Editing HOSTS.EQUIV file" window, enter any hosts and/or users that will need access to this execution host.
7. In the "Editing PBS MOM config file" window, accept the defaults.
8. Restart the execution host and log into it.
9. Stop the PBS MOM:

net stop pbs_mom

5.12.9 Install the New Version of PBS on the Server's Host

You must use the same password for the Administrator account on all hosts.

If you are installing PBS from the CD, put it in the CD-ROM drive, browse to your CD-ROM drive, and click on the PBS Professional icon.

If you are installing PBS from a download, save it to your hard drive and run the self-extracting `pbspro.exe` package, either in the same directory, or by giving the path to it.

Uninstall the old version of PBS:

1. Change your current directory so that it is not within `C:\INSTALLATION_BASE\PBS Pro`, and make sure there is no access occurring on any file in that hierarchy. Otherwise you will have to remove the hierarchy by hand.
2. Run the installation program by either clicking on the icon, or typing:

PBSpro_<version>-windows.exe

Under XP, SP2, you may see a warning saying, "The publisher could not be verified. Are you sure you want to run this software?" Ignore this message and click the "Run" button.

The installation package will ask whether you want to uninstall the previous version.

Answer *yes*.

If you see a popup window saying the hierarchy could not be removed, remove the hierarchy manually by going to a command window and typing the following. Do not use the `del` command.

rmdir /S /Q "C:\INSTALLATION_BASE\PBS Pro"

3. Reboot the server's host.

Install the new version of PBS:

1. Run the installation program by either clicking on the icon, or typing:

PBSpro_<version>-windows.exe

Under Windows XP SP2, you may see a warning saying, "The publisher could not be verified. Are you sure you want to run this software?" Ignore this message and click the "Run" button.

2. Check the installation location. You can change it, as long as the new location meets the requirements given in section 2.5, "Recommended PBS Configurations for Windows", on page 13.
3. Choose the "All" option.
4. When you are prompted for a license file location, enter the information.
5. Enter a non-empty password twice for the PBS service account. If you see "error 2245", check the password's length, complexity and history requirements. Look in the Active Directory guide or the Windows help page.
6. In the "Editing HOSTS.EQUIV file" window, enter any hosts and/or users that will need access to the server's host.
7. In the "PBS Server Nodes File" window, accept the defaults.
8. In the "PBS MOM Config File for local node" window, accept the defaults.
9. Restart the server's host, and log into it.
10. Stop the PBS MOM on the server's host:

```
net stop pbs_mom
```

5.12.10 Copy License and Resource Files

11. Copy the server's `resourcedef` file from the backup directory to the default directory. Type this command in a single line:

```
copy "%WINDIR%\TEMP\PBS Pro Backup\resourcedef.new" "INSTALLATION_BASE\PBS Pro\ home\server_priv\resourcedef"
```

5.12.11 Update Fairshare Definitions

This step is only necessary if you were not using `egroup:euser` as your fairshare entity, but wish to do so now. On the server/scheduler host:

1. Update `%WINDIR%\TEMP\PBS Pro Backup\sched_priv.work\resource_group`, with the new `egroup:euser` entities.
2. Restore the scheduler's modified directory:

```
copy "%WINDIR%\TEMP\PBS Pro Backup\sched_priv.work" "INSTALLATION_BASE\PBS
Pro\home\sched_priv\"
```

3. If you wish to keep your old fairshare resource usage data, use the `pbsfs` command to update the data.

Display the usage data. Both the old and the new entities will appear. Entities defined before the upgrade will appear in the unknown group. Entities defined in the previous step will be in the correct department but have no usage.

Display the usage data:

```
pbsfs -p
```

Assign the usage of the old entities to the new ones.

Be aware that the usage may not reflect what you want. For example, if your previous entities were `euser`, and now you have `egroup:euser`, and you have a user who is in more than one group, you will need to divide the user's usage so that it shows the usage by group that you want.

Old entity: Bob

```
% pbsfs -p
```

```
Bob : group: 1 cgrp: 5 shares: 10 Usage: 1005 Perc: 50%
```

The useful information is the name "Bob" and the usage of 1005

The newly defined entity is Math:Bob

```
Math:Bob 5 4 root 10
```

Assign Bob's usage to Math:Bob:

```
% pbsfs -s math:bob 1005
```

Remove the old entities:

```
% pbsfs -e
```

5.12.12 Start the New Server Without Defined Queues or Vnodes

When the new server starts it will have the default queue, workq, and server vnode already defined. You want to start the new server with empty configurations so that you can import your old settings.

1. Stop the new server:

```
qterm
```

2. Start the new server as a standalone, having it create a new database. Type this command in a single line:

```
"INSTALLATION_BASE\PBS Pro\exec\sbin\pbs_server" -N -t create
```

A message will appear saying "Create mode and server database exists, do you wish to continue?"

Type "y" to continue. The standalone server will exit after it creates the database.

Type `Ctrl-C` in order to get the terminal/command prompt back.

3. Start the new server as a Windows service:

```
net start pbs_server
```

4. Start the new scheduler as a Windows service:

```
net start pbs_sched
```

5.12.13 Set Server's License Location Attribute

Follow the steps in section 4.4.3.1, "Setting the Licensing Information in `pbs_license_info`", on page 87.

5.12.14 Replicate Queues, Server and Vnodes Configuration

1. Give the new server the old server's configuration, but modified for the new PBS. Type this command in a single line:

```
"INSTALLATION_BASE\PBS Pro\exec\bin\qmgr" < "%WINDIR%\TEMP\PBS Pro Backup\server.new"
```

2. List the queues in the new server:

```
"INSTALLATION_BASE\PBS Pro\exec\bin\qstat" -Q
```

3. Enable the queues in the new server. For each queue where its `enabled` attribute is `false`:

```
"INSTALLATION_BASE\PBS Pro\exec\bin\qenable" <queue name>
```

4. Replicate vnodes configuration, also modified for the new PBS. Type this command in a single line:

```
"INSTALLATION_BASE\PBS Pro\exec\bin\qmgr" < "%WINDIR%\TEMP\PBS Pro Backup\nodes.new"
```

5. Verify that the original configurations were read in properly. Type this command in a single line:

```
"INSTALLATION_BASE\PBS Pro\exec\bin\qmgr" -c "print server"  
"INSTALLATION_BASE\PBS Pro\exec\bin\pbsnodes" -a
```

5.12.15 Re-create Reservations

You must re-create each reservation that was on the old server, using the `pbs_rsub` command. Each reservation is created as a new reservation. You can use all of the information about the old reservation except for its start time. Be sure to give each reservation a start time in the future. Use the information stored in `%WINDIR%\TEMP\PBS Pro Backup\reservations`.

5.12.16 Change Port Specifications in `pbs.conf` for Old PBS

You must edit the `pbs.conf` file of the old PBS so that all old services use ports that won't clash with those of the new PBS. The old `pbs.conf` file was saved as `%WINDIR%\TEMP\PBS Pro Backup\pbs.conf`.

You must change the port numbers for all PBS related daemons: server, MOM, scheduler, and data service.

You edit the old configuration file so that the entries look like those in the following table:

Table 5-2: Entries in Old PBS Configuration File

New Entry in pbs.conf	Description
PBS_EXEC= %WINDIR%\TEMP\PBS Pro Backup\exec	This is the path for the old PBS_EXEC
PBS_HOME=%WINDIR%\TEMP\PBS Pro Backup\home	This is the path for the old PBS_HOME
PBS_START_SERVER=1	Unchanged
PBS_START_MOM=1	Unchanged
PBS_START_SCHED=1	Unchanged
PBS_SERVER=x17-64-rhel6	Unchanged
PBS_BATCH_SERVICE_PORT= 13001	This is the changed port number for the old server
PBS_DATA_SERVICE_PORT=13007	This is the changed port number for the old data service
PBS_MOM_SERVICE_PORT= 13002	This is the changed port number for the old MOM
PBS_SCHEDULER_SERVICE_PORT= 13003	This is the changed port number for the old scheduler

5.12.17 Start the Old Server

You must start the old server in order to move jobs to the new server. The old server must be started on alternate ports. Type the following commands without breaking the lines.

1. If you are upgrading from PBS Pro_5.3.3-wpl:

```
del "%WINDIR%\TEMP\PBS Pro Backup\home\server_priv\server.lock"
```

2. Tell PBS to use the `pbs.conf` file you saved in the backup directory, and to use the backup `exec` and `home` directories.

Do not put any double quotes in either path:

```
set PBS_CONF_FILE=%WINDIR%\TEMP\PBS Pro Backup\pbs.conf  
pbs-config-add "PBS_EXEC=%WINDIR%\TEMP\PBS Pro Backup\exec"  
pbs-config-add "PBS_HOME=%WINDIR%\TEMP\PBS Pro Backup\home"
```

3. Verify that the old server is using the `pbs.conf` saved in the backup directory:

```
echo %PBS_CONF_FILE%
```

4. Verify that `pbs.conf` contains the `exec` and `home` locations in the backup directory:

```
type %PBS_CONF_FILE%
```

5. Start the old server daemon in the same command prompt window as above, and assign these alternate port values:

```
"%WINDIR%\TEMP\PBS Pro Backup\exec\sbin\pbs_server" -N
```

For more information see section 6.2, “Starting and Stopping PBS: Windows”, on page 177 and “`pbs_server`” on page 90 of the PBS Professional Reference Guide.

After you start the old server, the command prompt will appear to hang, which is normal. Minimize this window and use a different command prompt. Closing the window or issuing `<Ctrl +C>` will stop the old server. If this happens, you will not be able to connect to the old server.

5.12.18 Set Configuration File Path to be Used by New Server

6. Tell PBS to use the `pbs.conf` file you saved in the backup directory, and to use the backup `exec` and `home` directories.

Do not put any double quotes in either path:

```
set PBS_CONF_FILE=INSTALLATION_BASE\PBS Pro\pbs.conf
```

5.12.19 Verify Old Server is Running on Alternate Ports

Verify that the old `pbs_server` is running on the alternate ports by going to another cmd prompt window and running, typed as a single line:

```
"%WINDIR%\TEMP\PBS Pro Backup\exec\bin\qstat" @<old server host>:13001
```

5.12.20 Migrate User Passwords From the Old Server to the New Server

You will want to migrate user passwords to the new server if possible. Passwords can only be migrated if both the old and new servers' `single_signon_password_enable` attributes are *True*. The following commands should be given in a single line:

1. Find out whether the old server's `single_signon_password_enable` attribute is *True*:

```
"%WINDIR%\TEMP\PBS Pro Backup\exec\bin\qmgr" -c "list s
single_signon_password_enable" <old server host>:13001
```

2. Find out whether the new server's `single_signon_password_enable` attribute is *True*:

```
"INSTALLATION_BASE\PBS Pro\exec\bin\qmgr" <new server host>:15001
Qmgr: list s single_signon_password_enable
```

3. If both attributes are *True*, you can migrate user passwords from the old server to the new server:

```
"INSTALLATION_BASE\PBS Pro\exec\bin\ pbs_migrate_users" <old server
host>:13001 <new server host>:15001
```

5.12.21 Move Existing Jobs to the New Server

You must move existing jobs from the old server to the new server. To do this, you run the new `qselect` and `qmove` commands, and give the new server's port number, 15001, in the destination. See "qmove" on page 169 of the PBS Professional Reference Guide. When moving jobs from reservation queues, be sure to move them into the equivalent new reservation queues.

There is one special case requiring an extra step. This is when the old server's `single_signon_password_enable` attribute is *False* and the new server's is *True*.

Give commands in a single line.

1. If the new server's `single_signon_password_enable` attribute is *True* and the old server's is *False*, temporarily set the new server's `single_signon_password_enable` to *False*:

```
"INSTALLATION_BASE\PBS Pro\exec\bin\qmgr" <new server host>:15001
Qmgr: set server single_signon_password_enable=false
```

2. You will need to verify later that all jobs have been moved. Print the list of jobs on the old server:

```
"%WINDIR%\TEMP\PBS Pro Backup\exec\bin\qstat" @<old server host>:13001
```

3. In another command prompt window, move each job in each queue.

This command may tie up the terminal. Create a file called `movejobs.bat`, containing the following lines. Replace `<old server host>` with the old server's host:

```
REM movejobs.bat
REM execute as follows:
REM
REM   movejobs <queue name>   (e.g. movejobs workq)
REM
setlocal ENABLEDELAYEDEXPANSION
for /F "usebackq" %%j in
    ( "INSTALLATION_BASE\PBS Pro\exec\bin\qselect" -q %1@<old server
    host>:13001^ )
do ( "INSTALLATION_BASE\PBS Pro\exec\bin\qmove"
    %1@<new server host>:15001
    %%j@<old server host>:13001
    )
```

Run `movejobs.bat` for each queue. Use the list of queue names saved when you disabled the queues.

For example, to move the jobs in `queue1` and `workq`, you would type:

```
cmd>movejobs queue1
cmd>movejobs workq
```

4. Verify that all jobs have been moved. Print the jobs on the new server:

```
"INSTALLATION_BASE\PBS Pro\exec\bin\qstat" @<new server host>:15001
```

5. Special case: If the old server's `single_signon_password_enable` attribute is *False* and the new server's was *True* (but is temporarily *False*), you must do the following

three steps:

- a. Apply a bad password hold to all the jobs on the new server. Create a file called `holdjobs.bat`, containing the following lines. Replace `<new server host>` with the new server's host:

```
REM pholdjobs.bat
REM execute as follows:
REM
REM   pholdjobs
REM
setlocal ENABLEDELAYEDEXPANSION
for /F "usebackq" %%j in
    ( ^"INSTALLATION_BASE\PBS Pro\exec\bin\
      qselect" -q @<new server host>:15001`
    do (
        "INSTALLATION_BASE\PBS
        Pro\exec\bin\qhold"
        -h p %%j@<new server host>:15001
    )
```

Run `pholdjobs.bat` by typing:

```
pholdjobs
```

- b. Change the new server's attribute back to *True*:

```
"INSTALLATION_BASE\PBS Pro\exec\bin\qmgr" <new server
host>:15001
Qmgr: set server single_signon_password_enable=true
```

- c. Each user with jobs on the new server must specify a password via `pbs_password`. See "pbs_password" on page 72 of the PBS Professional Reference Guide. Each user will type `pbs_password` and be prompted for a password.

5.12.22 Shut Down Old Server

1. Shut down the old server daemon:

```
"%WINDIR%\TEMP\PBS Pro Backup\exec\bin\qterm" -t quick <old server
host>:13001
```

5.12.23 Update New `sched_config`

Update the new scheduler's configuration file, in `INSTALLATION_BASE\PBS Pro\home\sched_priv\sched_config`, with any modifications that were made to the old `%WINDIR%\TEMP\PBS Pro Backup\sched_priv.work\sched_config`.

1. If you copied over your old scheduler log filter value, make sure that it has had 1024 added to it. If the value is less than 1024, add 1024 to it. For example, if the old log filter line is:

```
log_filter:256
```

change it to:

```
log_filter:1280
```

2. If it exists, replace the `strict_fifo` option with `strict_ordering`. If you do not, a warning will be printed in the log when the scheduler starts.

5.12.24 Start MOMs on Execution Hosts

1. On each execution host, start the MOM daemon as a Windows service:

```
net start pbs_mom
```

5.12.25 Optionally Start MOM on New Server's Host

If your old configuration had a MOM running on the server's host, and you wish to replicate the configuration, you can start a MOM on that machine.

Start the MOM daemon on the new server's host as a Windows service:

```
net start pbs_mom
```

5.12.26 Verify Communication Between Server and MOMs

Run `pbsnodes -a` on the server's host to see if it can communicate with the execution hosts in your complex. If a host is down, go to the problem host and restart the MOM:

```
net stop pbs_mom
```

```
net start pbs_mom
```

5.12.27 Enable Scheduling in the New Server

You must set the new server's scheduling attribute to *True* so that the scheduler will start jobs.

Enable scheduling in the new server:

```
"INSTALLATION_BASE\PBS Pro\exec\bin\qmgr" -c "set server scheduling=1"
```

5.12.28 Optionally Change PBS Service Account to Non-domain Administrator Account

If you want to run PBS Professional from a PBS service account which is not a domain administrator account, follow these steps:

1. Stop all PBS services:

```
net stop pbs_rshd
net stop pbs_mom
net stop pbs_server
net stop pbs_sched
```

2. Add the PBS service account to the local Administrators group:

```
net localgroup Administrators <domain name>\<name of PBS service account>
/add
```

3. Go to Active Directory, and make the PBS service account be only a member of the "Domain Users" group. Add the account to "Domain Users" on the local computer, and make it the primary group. Then remove the PBS service account's membership in the "Domain Admins" group.

4. Delegate explicit domain read privilege to the PBS service account. See section 2.5.2.4, "Delegating Read Access", on page 16.

5. Restart the PBS services:

```
net start pbs_rshd
net start pbs_mom
net start pbs_server
net start pbs_sched
```

Chapter 6

Starting and Stopping PBS

This chapter covers starting and stopping PBS.

6.1 Starting and Stopping PBS: UNIX and Linux

6.1.1 Methods for Starting Daemons

The PBS daemons can be started by two different methods. These methods are not equivalent. The first method is to use the PBS start/stop script, and the second is to run the command that starts the daemon. When you run the PBS start/stop script, PBS will create any vnode definition files. These are not created through the method of running the command that starts a daemon.

6.1.2 Recommendation to Start MOM First

We recommend starting MOMs before starting the server. This way, MOM will be ready to respond to the Server's "are you there?" ping, preventing the server from attempting to contact a MOM that is still down. This will cut down on inter-daemon traffic, especially in larger complexes.

6.1.3 The PBS Start/Stop Script

The PBS start/stop script is named `pbs_init.d`. To run it, you type the following:

```
<path to script>/pbs [start|stop|restart|status]
```

The script starts, stops, or restarts PBS daemons on the local machine. It can also be used to report the PID of any PBS daemon on the local machine. The PBS start/stop script reads the `pbs.conf` file to determine which components should be started.

The start/stop script runs at boot time, starting PBS upon bootup. You can run the script manually. The start/stop script runs on and affects only the local host.

When you run the PBS start/stop script, PBS will create any vnode definition files. These are not created through the method of running the command that starts a daemon.

The location of the PBS start/stop script depends on the OS, as shown in the following table:

Table 6-1: Location of Start/Stop Script

OS	Location of PBS Start/Stop Script
AIX	<code>/etc/rc.d/rc2.d/S90pbs</code>
Cray	If <code>/etc/init.d</code> exists: <code>/etc/init.d/pbs</code> Else: <code>/etc/rc.d/init.d/pbs</code>
HP-UX	<code>/sbin/init.d/pbs</code>
Linux	If <code>/etc/init.d</code> exists: <code>/etc/init.d/pbs</code> Else <code>/etc/rc.d/init.d/pbs</code>
Solaris	<code>/etc/init.d/pbs</code>

The command-line syntax for the start/stop script is:

```
<path to start/stop script>/pbs < status | stop | start | restart >
```

See “`pbs_init.d`” on page 51 of the PBS Professional Reference Guide.

6.1.3.1 Creation of Vnode Definition Files

When MOM is started via the PBS start/stop script, PBS creates any PBS reserved MOM configuration files, which contain vnode definitions. These are not created by the MOM itself, and will not be created when MOM is started using the `pbs_mom` command. Therefore, if you make changes to the number of CPUs or amount of memory that is available to PBS, or if a non-PBS process releases a `cpuset`, you should restart PBS in order to re-create the PBS reserved MOM configuration files. See section 3.4, “MOM and Vnode Configuration Files”, on page 41.

6.1.3.2 Effect of Start/Stop Script on Jobs

The PBS start script kills any running jobs when it is used.

6.1.3.3 Start/Stop Script Caveats

The PBS start/stop script uses the settings in `pbs.conf` to determine which daemons to start and stop. If you specify in `pbs.conf` that a daemon should not start, the script also will not stop it if it is running. Setting `PBS_START_MOM` to `0` effectively makes the start/stop script ignore the MOM. For example, if you do the following steps, the `pbs_mom` process is not stopped:

1. Start `pbs_mom`
2. Set `PBS_START_MOM` to `0`
3. Run the PBS start/stop script with `stop` as the argument

6.1.4 Starting Daemons Manually

6.1.4.1 Daemon Execution Requirements

The server, scheduler, MOM and the optional MOM Globus processes must run with the real and effective UID of root.

6.1.4.2 Required Permission

You must run `pbs_server`, `pbs_mom`, and `pbs_sched` as root.

6.1.4.3 Manually Starting MOM

The PBS MOM is started manually using the `pbs_mom` command. See “`pbs_mom`” on page 56 of the PBS Professional Reference Guide.

6.1.4.3.i Restarting MOM

You can restart MOM with the following options:

Table 6-2: MOM Restart Options

Option	Effect on Jobs
<code>pbs_mom</code>	Job processes will continue to run, but the jobs themselves are requeued.
<code>pbs_mom -r</code>	Processes associated with the job are killed. Running jobs are returned to the Server to be requeued or deleted. This option should not be used if the system has just been rebooted as the process numbers will be incorrect and a process not related to the job would be killed.
<code>pbs_mom -p</code>	Jobs which were running when MOM terminated remain running.

6.1.4.3.ii Preserving Existing Jobs When Restarting MOM

By default, when MOM is started, she allows running processes to continue to run, but tells the server to requeue her jobs. You can direct MOM to preserve running jobs and to track them, by using the `-p` option to the `pbs_mom` command. Start MOM with the command line:

```
PBS_EXEC/sbin/pbs_mom -p
```

Never use the `-p` option to `pbs_mom` after a reboot. See the next section.

6.1.4.3.iii Restarting MOM After a Reboot

When a UNIX/Linux operating system is first booted, it begins to assign process IDs (PIDs) to processes as they are created. PID 1 is always assigned to the system "init" process. As new ones are created, they are either assigned the next PID in sequence or the first empty PID found, which depends on the operating system implementation. Generally, the session ID of a session is the PID of the top process in the session.

The PBS MOM keeps track of the session IDs of the jobs. If only MOM is restarted on a system, those session IDs/PIDs have not changed and apply to the correct processes.

If the entire system is rebooted, the assignment of PIDs by the system will start over. Therefore the PID which MOM thinks belongs to an earlier job will now belong to a different later process. If you restart MOM with `-p`, she will believe the jobs are still valid jobs and the PIDs belong to those jobs. When she kills the processes she believes to belong to one of her earlier jobs, she will now be killing the wrong processes, those created much later but with the same PID as she recorded for that earlier job.

Never restart `pbs_mom` with the `-p` or the `-r` option following a reboot of the host system.

6.1.4.3.iv Killing Existing Jobs When Restarting MOM

If you wish to kill all existing processes, use the `-r` option to `pbs_mom`.

Start MOM with the command line:

```
PBS_EXEC/sbin/pbs_mom -r
```

6.1.4.3.v Starting MOM on the Altix

For a cpusetted Altix, MOM should be started using the PBS start/stop script.

6.1.4.3.vi Using Existing CPU and Memory for cpusets

By default, MOM removes existing cpusets when she starts. You can specify that MOM is to use existing CPU and memory allocations for cpusets by using the `-p` option to the `pbs_mom` command. This option also preserves running jobs. See “Options to `pbs_mom`” on page 58 of the PBS Professional Reference Guide.

Vnode definition files are not created when the `pbs_mom` command is used; use this only when you know that they are up to date already.

6.1.4.4 Manually Starting the Server

To start the server manually, do the following:

```
PBS_EXEC/sbin/pbs_server [options]
```

If, when the server was shut down, running jobs were killed and queued, then starting the server with the `-t hot` option puts those jobs back in the *Running* state first.

See “`pbs_server`” on page 90 of the PBS Professional Reference Guide for details and the options to the `pbs_server` command.

6.1.4.5 Manually Starting the Scheduler

To start the scheduler manually, do the following:

```
PBS_EXEC/sbin/pbs_sched [options]
```

There are no required options for the scheduler. See “`pbs_sched`” on page 88 of the PBS Professional Reference Guide for more information and a description of available options.

6.1.4.6 Running the Globus MOM

6.1.4.6.i Manually Starting Globus MOM

If you want support for Globus, start the Globus MOM at boot time. There are no required options. If you are starting the Globus MOM manually, run it with the line:

```
PBS_EXEC/sbin/pbs_mom_globus [options]
```

6.1.4.6.ii Restarting the Globus MOM

If Globus MOM is taken down and the host system continues to run, the Globus MOM should be restarted with the `-r` option. This directs Globus MOM to kill off processes running on behalf of a Globus job.

If the `pbs_mom_globus` process is restarted without the `-r` option, the assumption that will be made is that jobs have become disconnected from the Globus gatekeeper due to a system restart (cold start). Consequently, `pbs_mom_globus` will request that any Globus jobs that were being tracked and which were running be canceled and requeued.

See the **PBS Professional Programmer's Guide** or “`pbs_mom_globus`” on page 63 of the PBS Professional Reference Guide for a more complete explanation.

6.1.4.6.iii Globus MOM Caveats

The PBS start/stop script does not start the Globus MOM.

6.1.5 Stopping PBS

There are four ways to stop PBS. The first is to use the PBS start/stop script, the second is to use the `qterm` command, the third is to kill daemons by sending signals, and the fourth is to shut down the machine on which PBS is running.

6.1.5.1 Stopping PBS Using the Start/Stop Script

Stopping PBS using the start/stop script preserves running jobs.

When you use the PBS start/stop script, by typing “`pbs stop`” the following take place:

- The server gets a `qterm -t quick` (preserving jobs)
- MOM gets a SIGTERM: MOM terminates all running children and exits

6.1.5.2 Stopping PBS Using the `qterm` Command

The `qterm` command is used to shut down, selectively or inclusively, the various PBS components. If you have a failover server configured, then when the primary server is shut down, the secondary server becomes active unless you shut it down as well.

You can specify how running jobs are treated during shutdown by specifying the type of shutdown. The type of shutdown performed by the `qterm` command defaults to the following, which preserves running jobs:

```
qterm -t quick
```

The following command will shut down the primary server, the scheduler, and all MOMs in the complex. If configured, the secondary server will become active:

```
qterm -s -m
```

The following command will shut down the primary server, the secondary server, the scheduler, and all MOMs in the complex:

```
qterm -s -m -f
```

See “`qterm`” on page 226 of the PBS Professional Reference Guide.

6.1.5.2.i `qterm` Caveats

Shutting PBS down using the `qterm` command does not perform any of the other cleanup operations that are performed by the PBS start/stop script.

6.1.5.3 Stopping PBS by Sending Signals

The server does a quick shutdown, equivalent to receiving a `qterm -t quick`, upon receiving `SIGTERM`. See “`pbs_server`” on page 90 of the PBS Professional Reference Guide.

You can stop the scheduler by sending it `SIGTERM` or `SIGINT`. These result in an orderly shutdown of the scheduler.

The PBS MOM handles the following signals:

Table 6-3: Signals Handled by MOM

Signal	Effect
SIGTERM	If a MOM is killed with the signal SIGTERM, jobs are killed before MOM exits. Notification of the terminated jobs is not sent to the Server until the MOM is restarted. Jobs will still appear to be in the “R” (running) state.
SIGINT	If a MOM is killed with this signal, jobs are not killed before the MOM exits. MOM exits after cleanly closing network connections.
SIGKILL	If a MOM is killed with this signal, jobs are not killed before the MOM exits.

6.1.5.4 Shutting Down Host

When a host running PBS is shut down or rebooted, PBS is shut down via the start/stop script.

6.1.6 Stopping & Restarting One MOM

Should you ever need to stop a single MOM but leave jobs managed by her running, you have two options. The first is to send MOM a SIGINT. This will cause her to shut down in an orderly fashion. The second is to kill MOM with a SIGKILL (-9).

Note that MOM will need to be restarted with the `-p` option in order reattach to the jobs.

The PBS MOM can be stopped using the following signals:

Table 6-4: Signals to Stop MOM

Signal	Effect
SIGTERM	If a MOM is killed with the signal SIGTERM, jobs are killed before MOM exits. Notification of the terminated jobs is not sent to the Server until the MOM is restarted. Jobs will still appear to be in the “R” (running) state.
SIGINT	If a MOM is killed with this signal, jobs are not killed before the MOM exits. MOM exits after cleanly closing network connections.

Table 6-4: Signals to Stop MOM

Signal	Effect
SIGKILL	If a MOM is killed with this signal, jobs are not killed before the MOM exits.

You can restart MOM with the following options:

Table 6-5: MOM Restart Options

Option	Effect on Jobs
<code>pbs_mom</code>	Job processes will continue to run, but the jobs themselves are requeued.
<code>pbs_mom -r</code>	Processes associated with the job are killed. Running jobs are returned to the Server to be requeued or deleted. This option should not be used if the system has just been rebooted as the process numbers will be incorrect and a process not related to the job would be killed.
<code>pbs_mom -p</code>	Jobs which were running when MOM terminated remain running.

6.1.6.1 Stopping and Starting a cpuset MOM

Use the start/stop script to stop a MOM that is managing cpusets, when you are not preserving running jobs.

When you are preserving running jobs on a cpusetted machine, stop MOM using a signal, then restart MOM using the `-p` option to the `pbs_mom` command. For example:

```
kill -9 <MOM PID>
```

...

```
PBS_EXEC/sbin/pbs_mom -p
```

6.1.7 Impact of Shutdown / Restart on Running UNIX/Linux Jobs

The methods you can use to shut down PBS, and which daemons are shut down, will affect running jobs differently. You can leave normal jobs running during shutdown, but not job arrays. Any running subjobs of a job array are killed and requeued when the server is shut down.

The impact of a shutdown (and subsequent restart) on running jobs depends on the following:

- Whether you use the PBS start/stop script, a command, or a signal to shut down PBS
- How the Server (`pbs_server`) is shut down
- How MOM (`pbs_mom`) is shut down
- How MOM is restarted

6.1.7.1 Whether to Use Script, Command, or Signal for Shutdown and Restart

Use the `qterm` command to shut the server down when running jobs must be checkpointed before shutdown, allowed to run to completion before shutdown, or preserved through shutdown and restart. To preserve running jobs, use the `-p` option to the `pbs_mom` command when restarting MOM.

When the PBS start/stop script is used to stop PBS, MOM kills her jobs and exits. Jobs are requeued when MOM restarts.

6.1.7.2 Shutting Down Daemons Using Command

Choose one of the following recommended sequences, based on the desired impact on jobs, to stop and restart PBS:

- To allow running jobs to continue to run:

Shutdown:

```
qterm -t quick -m -s
```

Restart:

```
pbs_server -t warm
```

```
pbs_mom -p
```

```
pbs_sched
```

- To checkpoint and requeue checkpointable jobs, requeue rerunnable jobs, kill any non-

rerunnable jobs, then restart and run jobs that were previously running:

Shutdown:

```
qterm -t immediate -m -s
```

Restart:

```
pbs_mom
```

```
pbs_server -t hot
```

```
pbs_sched
```

- To checkpoint and requeue checkpointable jobs, requeue rerunnable jobs, kill any non-rerunnable jobs, then restart and run jobs without taking prior state into account:

Shutdown:

```
qterm -t immediate -m -s
```

Restart:

```
pbs_mom
```

```
pbs_server -t warm
```

```
pbs_sched
```

6.1.8 Checking Status of Daemons

You can check whether or not each daemon is running by using the PBS start script with the `status` option. To check the status of MOM, do the following on MOM's host:

```
<path to script>/pbs status
```

6.2 Starting and Stopping PBS: Windows

When PBS Professional is installed on Microsoft Windows, the PBS processes are registered as system services. As such, they will be automatically started and stopped when the system boots and shuts down. However, there may come a time when you need to manually stop or restart the PBS services (such as shutting them down prior to a PBS software upgrade).

6.2.1 Required Privilege

To stop or start PBS, you must have Administrator Privilege.

6.2.2 Methods for Stopping PBS on Windows

You can stop PBS on Windows using the `qterm` command or `net stop <service>`. For information on the `qterm` command, see “`qterm`” on page 226 of the PBS Professional Reference Guide.

The following example illustrates how to manually stop the PBS services. These lines must be typed at a Command Prompt.

```
net stop pbs_sched
net stop pbs_mom
net stop pbs_server
net stop pbs_rshd
```

6.2.3 Methods for Starting PBS on Windows

To start PBS on Windows, use the `net start <service>` command. You can specify startup options for the services using the method shown in the next section.

The following example shows how to manually start the PBS services. These lines must be typed at a Command Prompt.

```
net start pbs_server
net start pbs_mom
net start pbs_sched
net start pbs_rshd
```

6.2.3.1 Startup Options to PBS Services

You can use the startup options to the `pbs_server`, `pbs_sched`, and `pbs_mom` commands when starting the PBS services.

The procedure to specify startup options to the PBS Windows Services is as follows:

1. Go to *Start Menu->Control Panel->Performance and Maintenance->Administrative Tools->Services* (in Windows XP).
2. Select the PBS service you wish to alter. For example, if you select “*PBS_MOM*”, the MOM service dialog box comes up.
3. Enter the desired options in the “*Start parameters*” entry line. For example, to specify an alternate MOM configuration file, you might specify the following input:

On 32-bit Windows systems:

```
-c "\Program Files\PBS Pro\home\mom_priv\config2"
```

On 64-bit Windows systems:

```
-c "\Program Files (x86)\PBS Pro\home\mom_priv\config2"
```

4. Click on “*Start*” to start the specified service.

6.2.3.2 Startup Option Caveats

The Windows services dialog does not remember the “*Start parameters*” value when you close the dialog. You must specify the “*Start parameters*” value for each future restart.

6.2.4 Working with Version 2 Configuration Files on Windows

You can use the `pbs_mom` command to view and modify Version 2 configuration files, but you must use the command in standalone mode. For example, if you are inserting a new Version 2 configuration file, use the following:

```
pbs_mom -N -s insert <config file> <input script>
```

6.2.5 Running PBS Services in Standalone Mode

You can run the PBS services manually, in standalone mode and not as a Windows service, as follows:

```
Admin> pbs_server -N <options>
```

```
Admin> pbs_mom -N <options>
```

```
Admin> pbs_sched -N <options>
```

```
Admin> pbs_rshd -N <options>
```

6.2.6 Windows-specific Service Options

The PBS services have Windows-specific options:

`pbs_server -C`

The Server starts up, creates the database, and exits.

`-N`

(All services) The service runs in standalone mode, not as a Windows service.

6.2.7 Impact of Shutdown / Restart on Running Windows Jobs

The methods you can use to shut down PBS, and which daemons are shut down, will affect running jobs differently. You can leave normal jobs running during shutdown, but not job arrays. Any running subjobs of a job array are killed and requeued when the server is shut down.

The impact of a shutdown (and subsequent restart) on running jobs depends on whether you use `net stop` or the `qterm` command to shut down PBS, and how `pbs_mom` is restarted.

6.2.7.1 Choosing Shutdown and Restart Method

You can use either the `net stop` command or the `qterm` command to shut the server down.

When you use `net stop` to stop `pbs_server`, checkpointable jobs are checkpointed, killed and requeued, and all other jobs are killed and requeued. Jobs are not killed when `pbs_mom` is stopped via `net stop`; whether they are killed depends on how MOM is restarted.

Use the `qterm` command to shut the server down when running jobs must be checkpointed before shutdown, allowed to run to completion before shutdown, or preserved through shutdown and restart.

To preserve running jobs, use the `-p` option to the `pbs_mom` command when restarting MOM.

6.2.7.2 Shutting Down and Restarting Services

Choose one of the following recommended sequences, based on the desired impact on jobs, to stop and restart PBS.

- To allow running jobs to continue to run:

Shutdown:

```
qterm -t quick -m -s
```

Restart:

```
net start pbs_server (with -t warm startup option set)
```

```
net start pbs_mom (with -p startup option set)
```

```
net start pbs_sched
```

- To checkpoint and requeue checkpointable jobs, requeue rerunnable jobs, kill any non-rerunnable jobs, then restart and run jobs that were previously running:

Shutdown:

```
qterm -t immediate -m -s
```

or

```
net stop pbs_mom
```

```
net stop pbs_server
```

```
net stop pbs_sched
```

```
net stop pbs_rshd
```

Restart:

```
net start pbs_mom
```

```
net start pbs_server (with -t hot startup option set)
```

```
net start pbs_sched
```

- To checkpoint and requeue checkpointable jobs, requeue rerunnable jobs, kill any non-rerunnable jobs, then restart and run jobs without taking prior state into account:

Shutdown:

```
qterm -t immediate -m -s
```

Restart:

```
net start pbs_mom
```

```
net start pbs_server (with -t warm startup option set)
```

```
net start pbs_sched
```

6.3 When to Restart PBS

- If you make changes to the hardware or a change occurs in the number of CPUs or amount of memory that is available to PBS, such as a non-PBS process releasing a cpuset, you should restart PBS, by typing the following:
<path-to-script>/pbs restart
- After creating cpusets not managed by PBS on an Altix running ProPack 6 or 7 or Performance Suite 1. See section 11.7.1.1, "Manual Creation of cpusets Not Managed by PBS" on page 693 in the PBS Professional Administrator's Guide.
- After making changes to the `/etc/hosts` file. See section 2.1.3, "Name Resolution and Network Configuration", on page 7.
- After changing the name of the PBS service account.
- After changing the PBS service account to a non-domain administrator account

6.4 Caveats and Recommendations

6.4.1 Recommendation to Start MOM First

We recommend starting MOMs before starting the server. This way, MOM will be ready to respond to the Server's "are you there?" ping, preventing the server from attempting to contact a MOM that is still down. This will cut down on inter-daemon traffic, especially in larger complexes.

6.4.2 Recommendation to Offline Vnodes Before Stopping MOM

We recommend that you offline vnodes before stopping the MOM. The server tries to keep continual contact with each MOM. If you offline the vnode before stopping the MOM, the server does not try to stay in contact with the MOM. This reduces network traffic.

6.4.3 Start/Stop Script Caveats

The PBS start/stop script uses the settings in `pbs.conf` to determine which daemons to start and stop. If you specify in `pbs.conf` that a daemon should not start, the script also will not stop it if it is running. See section 6.1.3.3, "Start/Stop Script Caveats", on page 169.

6.4.4 Effect of Start/Stop Script on Jobs

The PBS start script kills any running jobs when it is used.

6.4.5 Starting MOM After Reboot

Never use the `-p` option to `pbs_mom` after a reboot. See section 6.1.4.3.iii, “Restarting MOM After a Reboot”, on page 170.

6.4.6 Globus MOM Caveats

The PBS start/stop script does not start the Globus MOM.

6.4.7 Caveats for Altix

For a cpusetted Altix, MOM must be started using the PBS start/stop script.

Appendix A: License Agreement

CAUTION!

PRIOR TO INSTALLATION OR USE OF THE SOFTWARE YOU MUST CONSENT TO THE FOLLOWING SOFTWARE LICENSE TERMS AND CONDITIONS BY CLICKING THE “I ACCEPT” BUTTON BELOW. YOUR ACCEPTANCE CREATES A BINDING LEGAL AGREEMENT BETWEEN YOU AND ALTAIR. IF YOU DO NOT HAVE THE AUTHORITY TO BIND YOUR ORGANIZATION TO THESE TERMS AND CONDITIONS, YOU MUST CLICK “I DO NOT ACCEPT” AND THEN HAVE AN AUTHORIZED PARTY IN THE ORGANIZATION THAT YOU REPRESENT ACCEPT THESE TERMS.

IF YOU, OR THE ORGANIZATION THAT YOU REPRESENT, HAS A MASTER SOFTWARE LICENSE AGREEMENT (“MASTER SLA”) ON FILE AT THE CORPORATE HEADQUARTERS OF ALTAIR ENGINEERING, INC. (“ALTAIR”), THE MASTER SLA TAKES PRECEDENCE OVER THESE TERMS AND SHALL GOVERN YOUR USE OF THE SOFTWARE.

MODIFICATION(S) OF THESE SOFTWARE LICENSE TERMS IS EXPRESSLY PROHIBITED. ANY ATTEMPTED MODIFICATION(S) WILL BE NONBINDING AND OF NO FORCE OR EFFECT UNLESS EXPRESSLY AGREED TO IN WRITING BY AN AUTHORIZED CORPORATE OFFICER OF ALTAIR. ANY DISPUTE RELATING TO THE VALIDITY OF AN ALLEGED MODIFICATION SHALL BE DETERMINED IN ALTAIR’S SOLE DISCRETION.

Altair Engineering, Inc. - Software License Agreement

THIS SOFTWARE LICENSE AGREEMENT, including any Additional Terms (together with the “Agreement”), shall be effective as of the date of YOUR acceptance of these software license terms and conditions (the “Effective Date”) and is between ALTAIR ENGINEERING, INC., 1820 E. Big Beaver Road, Troy, MI 48083-2031, USA, a Michigan corporation (“Altair”), and YOU, or the organization on whose behalf you have authority to accept these terms (the “Licensee”). Altair and Licensee, intending to be legally bound, hereby agree as follows:

1. DEFINITIONS. In addition to terms defined elsewhere in this Agreement, the following terms shall have the meanings defined below for purposes of this Agreement:

Additional Terms. Additional Terms are those terms and conditions which are determined by an Altair Subsidiary to meet local market conditions.

Documentation. Documentation provided by Altair or its resellers on any media for use with the Products.

Execute. To load Software into a computer's RAM or other primary memory for execution by the computer.

Global Zone: Software is licensed based on three Global Zones: the Americas, Europe and Asia-Pacific. When Licensee has Licensed Workstations located in multiple Global Zones, which are connected to a single License (Network) Server, a premium is applied to the standard Software License pricing for a single Global Zone.

ISV/Independent Software Vendor. A software company providing its products, (“ISV Software”) to Altair's Licensees through the Altair License Management System using Altair License Units.

License Log File. A computer file providing usage information on the Software as gathered by the Software.

License Management System. The license management system (LMS) that accompanies the Software and limits its use in accordance with this Agreement, and which includes a License Log File.

License (Network) Server. A network file server that Licensee owns or leases located on Licensee's premises and identified by machine serial number and/or HostID on the Order Form.

License Units. A parameter used by the LMS to determine usage of the Software permitted under this Agreement at any one time.

Licensed Workstations. Single-user computers located in the same Global Zone(s) that Licensee owns or leases that are connected to the License (Network) Server via local area network or Licensee's private wide-area network.

Maintenance Release. Any release of the Products made generally available by Altair to its Licensees with annual leases, or those with perpetual licenses who have an active maintenance agreement in effect, that corrects programming errors or makes other minor changes to the Software. The fees for maintenance and support services are included in the annual license fee but perpetual licenses require a separate fee.

Order Form. Altair's standard form in either hard copy or electronic format that contains the specific parameters (such as identifying Licensee's contracting office, License Fees, Software, Support, and License (Network) Servers) of the transaction governed by this Agreement.

Products. Products include Altair Software, ISV Software, and/or Suppliers' software; and Documentation related to all of the forgoing.

Proprietary Rights Notices. Patent, copyright, trademark or other proprietary rights notices applied to the Products, packaging or media.

Software. The Altair software identified in the Order Form and any Updates or Maintenance Releases.

Subsidiary. Subsidiary means any partnership, joint venture, corporation or other form of enterprise in which a party possesses, directly or indirectly, an ownership interest of fifty percent (50%) or greater, or managerial or operational control.

Suppliers. Any person, corporation or other legal entity which may provide software or documents which are included in the Software.

Support. The maintenance and support services provided by Altair pursuant to this Agreement.

Templates. Human readable ASCII files containing machine-interpretable commands for use with the Software.

Term. The term of licenses granted under this Agreement. Annual licenses shall have a 12-month term of use unless stated otherwise on the Order Form. Perpetual licenses shall have a term of twenty-five years. Maintenance agreements for perpetual licenses have a 12-month term.

Update. A new version of the Products made generally available by Altair to its Licensees that includes additional features or functionalities but is substantially the same computer code as the existing Products.

2. LICENSE GRANT. Subject to the terms and conditions set forth in this Agreement, Altair hereby grants Licensee, and Licensee hereby accepts, a limited, non-exclusive, non-transferable license to: a) install the Products on the License (Network) Server(s) identified on the Order Form for use only at the sites identified on the Order Form; b) execute the Products on Licensed Workstations in accordance with the LMS for use solely by Licensee's employees, or its onsite Contractors who have agreed to be bound by the terms of this Agreement, for Licensee's internal business use on Licensed Workstations within the Global Zone(s) as iden-

tified on the Order Form and for the term identified on the Order Form; c) make backup copies of the Products, provided that Altair's and its Suppliers' and ISV's Proprietary Rights Notices are reproduced on each such backup copy; d) freely modify and use Templates, and create interfaces to Licensee's proprietary software for internal use only using APIs provided that such modifications shall not be subject to Altair's warranties, indemnities, support or other Altair obligations under this Agreement; and e) copy and distribute Documentation inside Licensee's organization exclusively for use by Licensee's employees and its onsite Contractors who have agreed to be bound by the terms of this Agreement. A copy of the License Log File shall be made available to Altair automatically on no less than a monthly basis. In the event that Licensee uses a third party vendor for information technology (IT) support, the IT company shall be permitted to access the Software only upon its agreement to abide by the terms of this Agreement. Licensee shall indemnify, defend and hold harmless Altair for the actions of its IT vendor(s).

3. RESTRICTIONS ON USE. Notwithstanding the foregoing license grant, Licensee shall not do (or allow others to do) any of the following: a) install, use, copy, modify, merge, or transfer copies of the Products, except as expressly authorized in this Agreement; b) use any back-up copies of the Products for any purpose other than to replace the original copy provided by Altair in the event it is destroyed or damaged; c) disassemble, decompile or “unlock”, reverse translate, reverse engineer, or in any manner decode the Software or ISV Software for any reason; d) sublicense, sell, lend, assign, rent, distribute, publicly display or publicly perform the Products or Licensee's rights under this Agreement; e) allow use outside the Global Zone(s) or User Sites identified on the Order Form; f) allow third parties to access or use the Products such as through a service bureau, wide area network, Internet location or time-sharing arrangement except as expressly provided in Section 2(b); g) remove any Proprietary Rights Notices from the Products; h) disable or circumvent the LMS provided with the Products; or (i) link any software developed, tested or supported by Licensee or third parties to the Products (except for Licensee's own proprietary software solely for Licensee's internal use).

4. OWNERSHIP AND CONFIDENTIALITY. Licensee acknowledges that all applicable rights in patents, copyrights, trademarks, service marks, and trade secrets embodied in the Products are owned by Altair and/or its Suppliers or ISVs. Licensee further acknowledges that the Products, and all copies thereof, are and shall remain the sole and exclusive property of Altair and/or its Suppliers and ISVs. This Agreement is a license and not a sale of the Products. Altair retains all rights in the Products not expressly granted to Licensee herein. Licensee acknowledges that the Products are confidential and constitute valuable assets and trade secrets of Altair and/or its Suppliers and ISVs. Licensee agrees to take the same precautions necessary to protect and maintain the confidentiality of the Products as it does to protect its own information of a confidential nature but in any event, no less than a reasonable degree of care, and shall not disclose or make them available to any person or entity except as expressly provided in this Agreement. Licensee shall promptly notify Altair in the event any unauthorized person obtains access to the Products. If Licensee is required by any governmental

authority or court of law to disclose Altair's or its ISV's or its Suppliers' confidential information, then Licensee shall immediately notify Altair before making such disclosure so that Altair may seek a protective order or other appropriate relief. Licensee's obligations set forth in Section 3 and Section 4 of this Agreement shall survive termination of this Agreement for any reason. Altair's Suppliers and ISVs, as third party beneficiaries, shall be entitled to enforce the terms of this Agreement directly against Licensee as necessary to protect Supplier's intellectual property or other rights.

Altair and its resellers providing support and training to licensed end users of the Products shall keep confidential all Licensee information provided to Altair in order that Altair may provide Support and training to Licensee. Licensee information shall be used only for the purpose of assisting Licensee in its use of the licensed Products. Altair agrees to take the same precautions necessary to protect and maintain the confidentiality of the Licensee information as it does to protect its own information of a confidential nature but in any event, no less than a reasonable degree of care, and shall not disclose or make them available to any person or entity except as expressly provided in this Agreement.

5. MAINTENANCE AND SUPPORT. Maintenance. Altair will provide Licensee, at no additional charge for annual licenses and for a maintenance fee for paid-up licenses, with Maintenance Releases and Updates of the Products that are generally released by Altair during the term of the licenses granted under this Agreement, except that this shall not apply to any Term or Renewal Term for which full payment has not been received. Altair does not promise that there will be a certain number of Updates (or any Updates) during a particular year. If there is any question or dispute as to whether a particular release is a Maintenance Release, an Update or a new product, the categorization of the release as determined by Altair shall be final. Licensee agrees to install Maintenance Releases and Updates promptly after receipt from Altair. Maintenance Releases and Updates are subject to this Agreement. Altair shall only be obligated to provide support and maintenance for the most current release of the Software and the most recent prior release. **Support.** Altair will provide support via telephone and email to Licensee at the fees, if any, as listed on the Order Form. If Support has not been procured for any period of time for paid-up licenses, a reinstatement fee shall apply. Support consists of responses to questions from Licensee's personnel related to the use of the then-current and most recent prior release version of the Software. Licensee agrees to provide Altair with sufficient information to resolve technical issues as may be reasonably requested by Altair. Licensee agrees to the best of its abilities to read, comprehend and follow operating instructions and procedures as specified in, but not limited to, Altair's Documentation and other correspondence related to the Software, and to follow procedures and recommendations provided by Altair in an effort to correct problems. Licensee also agrees to notify Altair of a programming error, malfunction and other problems in accordance with Altair's then current problem reporting procedure. If Altair believes that a problem reported by Licensee may not be due to an error in the Software, Altair will so notify Licensee. Questions must be directed to Altair's specially designated telephone support numbers and email addresses. Support will also be available via email at Internet addresses designated by Altair. Support is available

Monday through Friday (excluding holidays) from 8:00 a.m. to 5:00 p.m. local time in the Global Zone where licensed, unless stated otherwise on the Order Form. **Exclusions.** Altair shall have no obligation to maintain or support (a) altered, damaged or Licensee-modified Software, or any portion of the Software incorporated with or into other software not provided by Altair; (b) any version of the Software other than the current version of the Software or the immediately prior release of the Software; (c) problems caused by Licensee's negligence, abuse or misapplication of Software other than as specified in the Documentation, or other causes beyond the reasonable control of Altair; or (d) Software installed on any hardware, operating system version or network environment that is not supported by Altair. Support also **excludes** configuration of hardware, non-Altair Software, and networking services; consulting services; general solution provider related services; and general computer system maintenance.

6. WARRANTY AND DISCLAIMER. Altair warrants for a period of ninety (90) days after Licensee initially receives the Software that the Software will perform under normal use substantially as described in then current Documentation. Supplier software included in the Software and ISV Software provided to Licensee shall be warranted as stated by the Supplier or the ISV. Copies of the Suppliers' and ISV's terms and conditions of warranty are available on the Altair Support website. Support services shall be provided in a workmanlike and professional manner, in accordance with the prevailing standard of care for consulting support engineers at the time and place the services are performed.

ALTAIR DOES NOT REPRESENT OR WARRANT THAT THE PRODUCTS WILL MEET LICENSEE'S REQUIREMENTS OR THAT THEIR OPERATION WILL BE UNINTERRUPTED OR ERROR-FREE, OR THAT IT WILL BE COMPATIBLE WITH ANY PARTICULAR HARDWARE OR SOFTWARE. ALTAIR EXCLUDES AND DISCLAIMS ALL EXPRESS AND IMPLIED WARRANTIES NOT STATED HEREIN, INCLUDING THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT. THE ENTIRE RISK FOR THE PERFORMANCE, NON-PERFORMANCE OR RESULTS OBTAINED FROM USE OF THE PRODUCTS RESTS WITH LICENSEE AND NOT ALTAIR. ALTAIR MAKES NO WARRANTIES WITH RESPECT TO THE ACCURACY, COMPLETENESS, FUNCTIONALITY, SAFETY, PERFORMANCE, OR ANY OTHER ASPECT OF ANY DESIGN, PROTOTYPE OR FINAL PRODUCT DEVELOPED BY LICENSEE USING THE PRODUCTS.

7. INDEMNITY. Altair will defend and indemnify, at its expense, any claim made against Licensee based on an allegation that the Software infringes a patent or copyright ("Claim"); provided, however, that this indemnification does not include claims which are based on Supplier software or ISV software, and that Licensee has not materially breached the terms of this Agreement, Licensee notifies Altair in writing within ten (10) days after Licensee first learns of the Claim; and Licensee cooperates fully in the defense of the claim. Altair shall have sole control over such defense; provided, however, that it may not enter into

any settlement binding upon Licensee without Licensee's consent, which shall not be unreasonably withheld. If a Claim is made, Altair may modify the Software to avoid the alleged infringement, provided however, that such modifications do not materially diminish the Software's functionality. If such modifications are not commercially reasonable or technically possible, Altair may terminate this Agreement and refund to Licensee the prorated license fee that Licensee paid for the then current Term. Perpetual licenses shall be pro-rated over a 36-month term. Altair shall have no obligation under this Section 7, however, if the alleged infringement arises from Altair's compliance with specifications or instructions prescribed by Licensee, modification of the Software by Licensee, use of the Software in combination with other software not provided by Altair and which use is not specifically described in the Documentation, and if Licensee is not using the most current version of the Software, if such alleged infringement would not have occurred except for such exclusions listed here. This section 7 states Altair's entire liability to Licensee in the event a Claim is made. No indemnification is made for Supplier and/or ISV Software.

8. LIMITATION OF REMEDIES AND LIABILITY. Licensee's exclusive remedy (and Altair's sole liability) for Software that does not meet the warranty set forth in Section 6 shall be, at Altair's option, either (i) to correct the nonconforming Software within a reasonable time so that it conforms to the warranty; or (ii) to terminate this Agreement and refund to Licensee the license fees that Licensee has paid for the then current Term for the nonconforming Software; provided, however that Licensee notifies Altair of the problem in writing within the applicable Warranty Period when the problem first occurs. Any corrected Software shall be warranted in accordance with Section 6 for ninety (90) days after delivery to Licensee. The warranties hereunder are void if the Software has been improperly installed, misused, or if Licensee has violated the terms of this Agreement.

Altair's entire liability for all claims arising under or related in any way to this Agreement (regardless of legal theory), shall be limited to direct damages, and shall not exceed, in the aggregate for all claims, the license and maintenance fees paid under this Agreement by Licensee in the 12 months prior to the claim on a prorated basis, except for claims under Section 7. **ALTAIR AND ITS SUPPLIERS AND ISVS SHALL NOT BE LIABLE TO LICENSEE OR ANYONE ELSE FOR INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES ARISING HEREUNDER (INCLUDING LOSS OF PROFITS OR DATA, DEFECTS IN DESIGN OR PRODUCTS CREATED USING THE SOFTWARE, OR ANY INJURY OR DAMAGE RESULTING FROM SUCH DEFECTS, SUFFERED BY LICENSEE OR ANY THIRD PARTY) EVEN IF ALTAIR OR ITS SUPPLIERS OR ITS ISVS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.** Licensee acknowledges that it is solely responsible for the adequacy and accuracy of the input of data, including the output generated from such data, and agrees to defend, indemnify, and hold harmless Altair and its Suppliers and ISVs from any and all claims, including reasonable attorney's fees, resulting from, or in connection with Licensee's use of the Software. No

action, regardless of form, arising out of the transactions under this Agreement may be brought by either party against the other more than two (2) years after the cause of action has accrued, except for actions related to unpaid fees.

9. UNITED STATES GOVERNMENT RESTRICTED RIGHTS. This section applies to all acquisitions of the Products by or for the United States government. By accepting delivery of the Products except as provided below, the government or the party procuring the Products under government funding, hereby agrees that the Products qualify as “commercial” computer software as that term is used in the acquisition regulations applicable to this procurement and that the government's use and disclosure of the Products is controlled by the terms and conditions of this Agreement to the maximum extent possible. This Agreement supersedes any contrary terms or conditions in any statement of work, contract, or other document that are not required by statute or regulation. If any provision of this Agreement is unacceptable to the government, Vendor may be contacted at Altair Engineering, Inc., 1820 E. Big Beaver Road, Troy, MI 48083-2031; telephone (248) 614-2400. If any provision of this Agreement violates applicable federal law or does not meet the government's actual, minimum needs, the government agrees to return the Products for a full refund.

For procurements governed by DFARS Part 227.72 (OCT 1998), the Software, except as described below, is provided with only those rights specified in this Agreement in accordance with the Rights in Commercial Computer Software or Commercial Computer Software Documentation policy at DFARS 227.7202-3(a) (OCT 1998). For procurements other than for the Department of Defense, use, reproduction, or disclosure of the Software is subject to the restrictions set forth in this Agreement and in the Commercial Computer Software - Restricted Rights FAR clause 52.227-19 (June 1987) and any restrictions in successor regulations thereto.

Portions of Altair's PBS Professional Software and Documentation are provided with RESTRICTED RIGHTS. Use, duplication, or disclosure by the Government is subject to restrictions as set forth in subdivision(c)(1)(ii) of the rights in the Technical Data and Computer Software clause in DFARS 252.227-7013, or in subdivision (c)(1) and (2) of the Commercial Computer Software-Restricted Rights clause at 48 CFR52.227-19, as applicable.

10. CHOICE OF LAW AND VENUE. This Agreement shall be governed by and construed under the laws of the state of Michigan, without regard to that state's conflict of laws principles except if the state of Michigan adopts the Uniform Computer Information Transactions Act drafted by the National Conference of Commissioners of Uniform State Laws as revised or amended as of June 30, 2002 (“UCITA”) which is specifically excluded. This Agreement shall not be governed by the United Nations Convention on Contracts for the International Sale of Goods, the application of which is expressly excluded. Each Party waives its right to a jury trial in the event of any dispute arising under or relating to this Agreement. Each party agrees that money damages may not be an adequate remedy for breach

of the provisions of this Agreement, and in the event of such breach, the aggrieved party shall be entitled to seek specific performance and/or injunctive relief (without posting a bond or other security) in order to enforce or prevent any violation of this Agreement.

11. [RESERVED]

12. Notice. All notices given by one party to the other under the Agreement or these Additional Terms shall be sent by certified mail, return receipt requested, or by overnight courier, to the respective addresses set forth in this Agreement or to such other address either party has specified in writing to the other. All notices shall be deemed given upon actual receipt.

Written notice shall be made to:

Altair: **Licensee Name & Address:**
Altair Engineering, Inc. _____
1820 E. Big Beaver Rd _____
Troy, MI 48083 _____
Attn: Tom M. Perring Attn: _____

13. TERM. For annual licenses, or Support provided for perpetual licenses, renewal shall be automatic for each successive year (“Renewal Term”), upon mutual written execution of a new Order Form. All charges and fees for each Renewal Term shall be set forth in the Order Form executed for each Renewal Term. All Software licenses procured by Licensee may be made coterminous at the written request of Licensee and the consent of Altair.

14. TERMINATION. Either party may terminate this Agreement upon thirty (30) days prior written notice upon the occurrence of a default or material breach by the other party of its obligations under this Agreement (except for a breach by Altair of the warranty set forth in Section 8 for which a remedy is provided under Section 10; or a breach by Licensee of Section 5 or Section 6 for which no cure period is provided and Altair may terminate this Agreement immediately) if such default or breach continues for more than thirty (30) days after receipt of such notice. Upon termination of this Agreement, Licensee must cease using the Software and, at Altair's option, return all copies to Altair, or certify it has destroyed all such copies of the Software and Documentation.

15. GENERAL PROVISIONS. Export Controls. Licensee acknowledges that the Products may be subject to the export control laws and regulations of the United States and other countries, and any amendments thereof. Licensee agrees that Licensee will not directly or indirectly export the Products into any country or use the Products in any manner except in compliance with all applicable U.S. and other countries export laws and regulations. **Notice.** All notices given by one party to the other under this Agreement shall be sent by certified mail, return receipt requested, or by overnight courier, to the respective addresses set forth in

this Agreement or to such other address either party has specified in writing to the other. All notices shall be deemed given upon actual receipt. **Assignment.** Neither party shall assign this Agreement without the prior written consent of other party, which shall not be unreasonably withheld. All terms and conditions of this Agreement shall be binding upon and inure to the benefit of the parties hereto and their respective successors and permitted assigns. **Waiver.** The failure of a party to enforce at any time any of the provisions of this Agreement shall not be construed to be a waiver of the right of the party thereafter to enforce any such provisions. **Severability.** If any provision of this Agreement is found void and unenforceable, such provision shall be interpreted so as to best accomplish the intent of the parties within the limits of applicable law, and all remaining provisions shall continue to be valid and enforceable. **Headings.** The section headings contained in this Agreement are for convenience only and shall not be of any effect in constructing the meanings of the Sections. **Modification.** No change or modification of this Agreement will be valid unless it is in writing and is signed by a duly authorized representative of each party. **Conflict.** In the event of any conflict between the terms of this Agreement and any terms and conditions on a Licensee Purchase Order or comparable document, the terms of this Agreement shall prevail. Moreover, each party agrees any additional terms on any Purchase Order or comparable document other than the transaction items of (a) item(s) ordered; (b) pricing; (c) quantity; (d) delivery instructions and (e) invoicing directions, are not binding on the parties. In the event of a conflict between the terms of this Agreement, and the Additional Terms, the Agreement shall take precedence. **Entire Agreement.** This Agreement, the Additional Terms, and the Order Form(s) attached hereto constitute the entire understanding between the parties related to the subject matter hereto, and supersedes all proposals or prior agreements, whether written or oral, and all other communications between the parties with respect to such subject matter. This Agreement may be executed in one or more counterparts, all of which together shall constitute one and the same instrument. **Execution.** Copies of this Agreement executed via original signatures, facsimile or email shall be deemed binding on the parties.

Index

Numerics

32-bit 102

64-bit 102

A

account

 PBS service 14

Active Directory 13, 18

Admin 13

Administrators 13

AIX 33

API

 SGI ProPack 8

authorization 12, 20

Avail_Global 85

Avail_Local 85

B

backup directory

 overlay upgrade 110, 114

 Windows upgrade 148

bad password hold

 Windows upgrade 164

Boolean Resources

 migration under UNIX 122

 Windows upgrade 151

C

Cannot delete busy object

 migration under UNIX 130

CD-ROM 27

client commands 3

cluster 11

commands 3

Comprehensive System Accounting 9

D

Delegation 13

deprecated terms

 migration 123

 Windows upgrade 152

Displaying Licensing Information 91

DNS 66

Domain Admin Account 14

Domain Admins 14

Domain User Account 14

Domain Users 14

domains

 mixed 20

E

empty queue, node configurations

 migration under UNIX 127

Enterprise Admins 14

Index

- error 2245
 - Windows upgrade 154
- Execution Hosts
 - installing on during Windows upgrade 152
- executor 3
- F**
- Failover 90
- failover
 - migration 116, 125
- File
 - .rhosts 12, 21
 - .shosts 12
 - hosts.equiv 70
- file permission
 - Windows upgrade 150
- Files
 - .rhosts 73
 - host.equiv 12, 21
 - hosts.equiv 21, 67, 72, 73
 - PBS Startup Script 168
 - pbs.conf 71
 - rhosts 12, 21
 - services 31
- FLicenses 85
- Floating License 83
- G**
- gethostbyaddr 30
- Globus
 - gatekeeper 172
 - MOM, starting 172
- H**
- h flag 112, 117
- hierarchy could not be removed
 - Windows upgrade 155
- High_Use 85
- hosts.equiv 73
- I**
- IETF 7, 30
- Install Account 14
- installation
 - cluster 29
 - Unix/Linux 34
 - Windows 2000 65
- J**
- Job
 - Executor (MOM) 3
 - Scheduler 3
- L**
- License Server 83
- license server configuration
 - redundant 83
- License Server List Configuration 83
- license_count 85
- Licensing 84, 90
- licensing 90
- Licensing Errors 98
- Linux job container 9
- log_filter
 - update for upgrade 131
- M**
- manager 3
 - commands 3
- Memory-only Vnode 83
- migrating user passwords
 - Windows upgrade 162
- migration upgrade
 - deprecated terms 123
 - UNIX 119
 - Windows 145
- mixed domains 20
- MOM 3, 11
- monitoring 1
- Move Existing Jobs to New Server 130
 - Windows upgrade 162

Index

- movejobs.bat
 - Windows upgrade 163
- moving jobs
 - migration upgrade under UNIX 129
- N**
- network
 - addresses 30
 - ports 30
 - services 31
- NFS 8
- node attributes
 - saving during Windows upgrade 148
- node-level resources
 - flags 112, 117
- nodes configuration file
 - migration upgrade under UNIX 122
- NQS 1
- NTFS 17, 19
- O**
- operator
 - commands 3
- output files 12
- overlay upgrade 101
 - backup directory 110, 114
 - pbs.conf 111
 - scheduler configuration 112, 117
 - UNIX and Linux 109
- P**
- package installer 35
- package, installing
 - migration under UNIX 126
- password
 - single-signon 21
 - Windows 18, 20
- PBS service account 14
- PBS Startup Script 168
- pbs.conf 32, 44
- PBS_EXEC 32, 71
- PBS_EXEC/pbs_sched_config
 - overlay upgrade 112, 117
- PBS_HOME 32, 71
- pbs_license_file_location 85
- pbs_license_linger_time 86
- pbs_license_max 86
- pbs_license_min 86
- pbs_migrate_users
 - Windows upgrade 162
- pbs_mom 3, 11, 31, 66
 - starting during overlay 113
- pbs_password
 - Windows upgrade 164
- pbs_probe 79
- pbs_rshd 72
 - stopping during upgrade 149, 166
- pbs_sched 3, 10, 11, 31, 66
- pbs_server 2, 10, 11, 31, 66
- pbs_server -t create
 - migration under UNIX 127
- pbs-config-add
 - Windows upgrade 161
- pbsnodes 70
- pholdjobs.bat
 - Windows upgrade 164
- POSIX 3
- PROP type=boolean flag=h
 - migration 122
- Q**
- qalter 22
- qmgr 78
- qsub 22
- queuing 1
- Quick Start Guide v
- R**
- rcp 8
- redundant license server configuration 83
- redundant license servers 90
- Release Notes

Index

- upgrade recommendations 101
- resourcedef
 - migration under UNIX 120
- resources_available
 - migration 122

S

- sched_config
 - update during upgrade 131
 - updating for Windows upgrade 165
- Scheduler 3, 11
 - policies 1
- scheduler configuration File
 - overlay upgrade 112, 117
- scheduler log filter
 - overlay upgrade 112, 117
- Schema Admins 14
- scp 8, 12
- Secure Copy 12
- Server 2, 11
- server attributes
 - saving during Windows upgrade 148
- server's configuration
 - migration upgrade under UNIX 120
- Server's Host
 - installing on during Windows upgrade 154
- server's node attributes
 - migration under UNIX 120
- service account
 - PBS 14
- SGI
 - Altix 8
 - ProPack Library 8
- SGI cpusets 8
- single_signon_password_enable 162
 - moving jobs 162, 163
 - Windows upgrade 162
- single-signon 21
- ssh 12
- Starting
 - MOM 169

- PBS 167
- Schedulr 171
- Server 171
- Startup Script 168
- system daemons 2

T

- tarfile
 - migration under UNIX 121
 - overlay upgrade 110, 114
- Three-server Configuration 83

U

- update resources
 - Windows upgrade 151
- upgrade
 - migration 101
 - migration under UNIX 119
 - migration under Windows 145
 - overlay 101
- upgrade under Windows
 - migrating user passwords 162
- Upgrading 97
- upgrading
 - UNIX and Linux 107
 - Windows 145
- Used 85
- user
 - commands 2, 3
- User Guide v
- User's Guide 73

W

- Windows 20, 21
 - password 21
- Windows 2000 66
- Windows Configuration in a Standalone Environment 17
- Windows XP 66
- workload management 1

Index

X

X forwarding 79

xauth 79

Index
