

PBS Professional® 14.2.3

Release Notes



 Altair | PBS Works™

PBS Works is a division of  Altair

PBS Professional 14.2.3 Release Notes, updated 7/19/17.

Please send any questions or suggestions for improvements to agu@altair.com.

Copyright © 2003-2017 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.

Terms of use for this software are available online at

<http://www.pbspro.com/UserArea/agreement.html>

This document is proprietary information of Altair Engineering, Inc.

Please send any questions or suggestions for improvements to agu@altair.com.

Contact Us

For the most recent information, go to the PBS Works website, www.pbsworks.com, select "My PBS", and log in with your site ID and password.

Altair

Altair Engineering, Inc., 1820 E. Big Beaver Road, Troy, MI 48083-2031 USA www.pbsworks.com

Sales

pbssales@altair.com 248.614.2400

Technical Support

Need technical support? We are available from 8am to 5pm local times:

Location	Telephone	e-mail
Australia	+1 800 174 396	anz-pbssupport@india.altair.com
China	+86 (0)21 6117 1666	es@altair.com.cn
France	+33 (0)1 4133 0992	pbssupport@europe.altair.com
Germany	+49 (0)7031 6208 22	pbssupport@europe.altair.com
India	+91 80 66 29 4500 +1 800 425 0234 (Toll Free)	pbs-support@india.altair.com
Italy	+39 800 905595	pbssupport@europe.altair.com
Japan	+81 3 5396 2881	pbs@altairjp.co.jp
Korea	+82 70 4050 9200	support@altair.co.kr
Malaysia	+91 80 66 29 4500 +1 800 425 0234 (Toll Free)	pbs-support@india.altair.com
North America	+1 248 614 2425	pbssupport@altair.com
Russia	+49 7031 6208 22	pbssupport@europe.altair.com
Scandinavia	+46 (0)46 460 2828	pbssupport@europe.altair.com
Singapore	+91 80 66 29 4500 +1 800 425 0234 (Toll Free)	pbs-support@india.altair.com
South America	+55 11 3884 0414	br_support@altair.com
UK	+44 (0)1926 468 600	pbssupport@europe.altair.com

Contents

1	Supported Platforms and New Features	1
1.1	Supported Platforms	1
1.2	MPI	1
1.3	MUNGE	2
1.4	CSA	2
1.5	Altair License Server	2
1.6	New Features	3
2	Special Notes	5
2.1	Upgrading	5
2.2	Licensing	7
2.3	Installation	8
2.4	Caveats for Uninstalling on Linux	8
2.5	Libraries	8
2.6	General	9
2.7	External Interface Changes	11
2.8	Documentation	17
2.9	Hooks	18
2.10	SGI	20
2.11	MPI	21
2.12	Linux	21
3	Bugs	23
3.1	Modifications and Bug Fixes in PBS 14.2.3	23
3.2	Modifications and Bug Fixes in PBS 14.2.2	23
3.3	Known Bugs, Errors, and Warnings	24
4	Legal Notices	33
4.1	avltree	33
4.2	popen.c	33
4.3	Cray	34
4.4	libical	34
4.5	Expat	34
4.6	hwloc	35
4.7	editline	35
4.8	MinGwEditLine	36
	Index	37

Table of Contents

Supported Platforms and New Features

1.1 Supported Platforms

For the following platforms, you can download all PBS components:

Table 1-1: Supported Platforms

Maker	Version	Chipset	Notes
Red Hat Enterprise Linux	(RHEL) 6 and 7	x86_64	
CentOS	6.x and 7.x	x86_64	
SGI	ICE, UV and cluster products with SGI Performance Suite 1	x86_64	
SuSE	SLES 11 and 12	x86_64	

1.1.1 Using PBS on Cray

For Cray: Use PBS 13.0.40x; do not use PBS 14.2.3.

1.1.2 Using PBS on SGI with Power Awareness

For SGI with Power Awareness: Use PBS 13.0.50x; do not use PBS 14.2.3.

1.1.3 Using PBS on Windows

For Windows, use PBS 13.1; do not use PBS 14.2.3.

1.2 MPI

PBS runs jobs with all known implementations of MPI. PBS is integrated more tightly with some versions, providing support for resource tracking and job cleanup after a failure. These are listed here:

Table 1-2: Integrated MPIs

MPI	Version	Platform	Notes
MPICH	1.2.5, 1.2.6, and 1.2.7	Linux	Support is deprecated as of 14.2.1

Table 1-2: Integrated MPIs

MPI	Version	Platform	Notes
MPICH2	1.0.3, 1.0.5, 1.0.7	Linux	
MPICH-GM	mpich-1.2.6.14b	Linux	Support is depre- recated as of 14.2.1
Intel MPI	2.0.022, 3, and 4.0.3	Linux	
HP MPI	2.0.0	Linux	
Platform MPI	8.0	Linux	
SGI MPI (MPT)		Linux on SGI platforms, including over InfiniBand	
MVAPICH	1.2.7	Linux	Support is depre- recated as of 14.2.1
MVAPICH2	1.7	Linux	

The following MPIs have taken the steps to integrate tightly with PBS:

Table 1-3: MPIs Self-integrated with PBS

MPI	Version	Platform	
LAM/MPI	6.5.9, 7.0.6, 7.1.1	Linux	
OpenMPI	1.4.2		

1.3 MUNGE

PBS can now be configured to use MUNGE version 0.5.11

1.4 CSA

Support for CSA is deprecated.

Any system running PBS and Performance Suite will not support CSA.

1.5 Altair License Server

For a list of supported platforms for the Altair License Server, see the License Server release notes. These are available from the PBS Works Documentation Download Area:

<https://secure.altair.com/UserArea/docs.php>

1.6 New Features

Table 1-4: New Features in PBS 14.2.2

JIRA Issue	Summary
PBS-17523	Upgrade to Python 2.7.13
PBS-17587	Add Delegate=yes to pbs.service file

2

Special Notes

2.1 Upgrading

2.1.1 Choosing Between Overlay and Migration Upgrades

We recommend a migration upgrade for the following:

- On Cray: when upgrading from pre-11.0 to 11.0 or newer
- If you are running a PBS version older than 5.3
- When moving between 32-bit and 64-bit versions of PBS
- When certain European or Japanese characters are stored in the datastore on versions of PBS older than 12.0

We recommend an overlay upgrade for the following:

- On Cray: when upgrading from 11.0 and newer
- Other Linux/UNIX platforms: when upgrading from 5.3 or newer

See the PBS Professional Installation & Upgrade Guide.

2.1.2 Order of Upgrades

When upgrading a PBS complex, make sure that all MoMs are upgraded before upgrading the server. Do not start the new server until all MoMs have been upgraded and restarted.

2.1.3 User Performing Upgrade Must be Database Creator

In order to upgrade the database, the user performing the upgrade must be the same user who created the database. If you will use a different database user after the upgrade, make that change after the upgrade. Perform the upgrade as the user who created the database, then change the user to the new desired user. Follow these steps:

1. Run this command on the PBS database in order to identify the user who originally created the database (this is one line):

```
su - <data_user> -s /bin/sh -c "LD_LIBRARY_PATH=$PBS_EXEC/pgsql/lib $PBS_EXEC/pgsql/bin/psql -U
<data_user> -p <data_port> -d pbs_datastore"
```

data_port defaults to "15007"

data_user defaults to "pbsdata"

2. Run this query to get the list of users in the database:

```
pbs_datastore=# select pg_authid.oid, rolname from pg_authid;
```

```
oid | rolname
-----+-----
 10 | pbsdata
16541 | ashr
```

(2 rows)

3. Find the original user who created the database (this is one line):

```
pbs_datastore=# select pg_authid.oid, rolname from pg_authid where pg_authid.oid=10;
```

```
oid | rolname
-----+-----
 10 | pbsdata
```

(1 row)

4. Exit from the PBS database
5. If the user does not exist, create the user in the system
6. Update the current database user to the original user who created the database:

```
pbs_ds_password -C pbsdata
```

7. Perform an overlay upgrade
8. Reset the current database user to the desired username:

```
pbs_ds_password -C "username"
```

2.1.4 Rolling upgrades

Rolling upgrades are not supported. However, you can perform a rolling upgrade from 13.0 to a newer release (you cannot go from pre-13.0 to post-13.0), and in the second paragraph, we provide some information about how to do a rolling upgrade. You should upgrade all MoMs and the server/scheduler at the same time; an older MoM cannot interoperate with a newer server. For example, a 14.2.3 server will report all pre-13.0 MoMs as down.

Before performing a rolling upgrade, consider stopping the scheduler before beginning the upgrade. When MoMs are upgraded to 14.2.3, they become unable to communicate with the server and vice versa. The scheduler cannot see or use the vnodes, and the MoMs cannot send obits to the server. After the server is upgraded, communication is restored, the scheduler can send jobs, and MoMs can send obits to the server. However, if the scheduler is running during a rolling upgrade, two things may happen:

- a. Jobs that were top jobs may lose their top job status, and therefore any resources the scheduler has designated for them may be assigned to other jobs. This can delay top jobs.
- b. As vnodes become unavailable to the scheduler, users may become concerned because their jobs may stay queued with a comment saying that their job can never run. Once all vnodes and the server are upgraded, the jobs can run.

2.1.5 Altair License Server

PBS 11.0 and newer uses the Altair License Server. If you are not using socket licenses for all machines, you should install and configure the v13.0 or newer Altair License Server before you upgrade PBS. See the PBS Professional Installation & Upgrade Guide and the Altair License Manager Installation and Operations Guide.

2.1.6 Upgrading Caveats

- Warning when upgrading from 14.2.1: avoid using “`rpm -U`”. Instead, uninstall 14.2.1 using “`rpm -e`”, then install the next version using “`rpm -i`”.
- The `qrerun` command does not requeue job arrays. Instead, it leaves them in the *B* (begun) state. During a migration upgrade, you can work around this problem by using the `qhold` command on the job array, then using `qmove` to move it to the new server, then using the `qrls` command.

2.2 Licensing

- As of 11.3, PBS can license jobs using socket licenses and/or floating licenses. Floating licenses are managed via the Altair License Server. You CANNOT use a FLEX server with PBS 13.1. Follow the instructions in the Altair License Manager Installation and Operations Guide.
- The Licensing chapter in the PBS Professional Installation & Upgrade Guide describes how to configure PBS Professional to use the Altair License Server and/or socket licenses.
- If you find that the PBS server host becomes unresponsive a short time after startup, the server may be trying to contact the wrong license server. See the Troubleshooting chapter in the PBS Professional Administrator's Guide.
- The `pbs_license_file_location` server attribute is deprecated. The installation script unsets this attribute during an upgrade. This is intentional. Use instead the `pbs_license_info` server attribute.
- The minimum number of PBS Works licenses that can be checked out is one (1). If the `pbs_license_min` server attribute is set to zero or unset, PBS automatically sets the number to one.
- Starting with version 13.1, the default value for the `pbs_license_linger_time` server attribute is set to 1 year (31536000 seconds). The previous default value was 3600 seconds. This means that any unused licenses are

checked back in to the license server after a year, instead of one hour. This value can be overridden by setting the value of `pbs_license_linger_time` via `qmgr`.

- If you are using socket-based licenses to run PBS Professional, you will be asked for the host ID when generating the license file. You will need to run the `almutil` utility to gather this information. You can get the `almutil` utility from either:
 - An existing LM-X Altair License Server installation
 - A tar file containing `almutil` binaries available on the Altair Registered User License Manager page after selecting "Create New License". Extract the version you need for your PBS server platform.
- You can refer to the Altair License Manager Installation and Operations Guide (also available from the Registered User Area site) if you are unsure which binary you should use for your particular platform.

2.3 Installation

- Any `PBS_START_*` parameters set in the environment are not picked up and set in `pbs.conf`. You must specify these in `pbs.conf`; do not export them.
- Create data service account:

Under UNIX/Linux, we recommend that the PBS administrator creates a new data service account called "pbsdata" before installing PBS. This account must be created correctly. See the PBS Professional Installation & Upgrade Guide for details.

2.4 Caveats for Uninstalling on Linux

Using `rpm -e`, even on an older package than the one you are currently using, will cause any currently running PBS daemons to shut down, and will also remove the system V init and/or systemd service startup files. If you wish to remove an older rpm without these effects, use `rpm -e --noscripts`.

2.5 Libraries

2.5.1 Dependency on libdl

If you are integrating for example OpenMPI with the PBS tm API by using the `--with-tm` option, you need to use `libdl`. Before you configure and make OpenMPI, you can do this:

```
export LIBS=-ldl
```

As of 13.1, `libpbs.a` uses `dlopen`, `dlclose`, `dlsym` etc. for MUNGE-related APIs etc.

2.6 General

- SLES 12 SP2 has a newly introduced limit on the number of threads/processes. The version of `systemd` shipped with SLES 12 SP2 contains support for the PIDs cgroup controller:

https://www.suse.com/releasenotes/x86_64/SUSE-SLES/12-SP2/#fate-320358

Under the new `systemd` version, each init script or `systemd` service is limited to 512 threads/processes by default. This could cause issues if a MOM daemon runs many simultaneous jobs, or a compute node launcher daemon attempts to spawn many processes for a job.

Here's how to increase the limit beyond the default:

- If using a `systemd` service file, add `TasksMax=N` to the [Service] section. *N* can be a specific number, or special value infinity.
- If using an init script, create the file `/etc/systemd/system/<init script name>.service.d/override.conf` with these contents:

```
[Service]
TasksMax=N
```

Earlier versions of `systemd` that don't support the PIDs cgroup controller simply ignore the `TasksMax` setting.

- In the server `periodic` hook that runs `pbs_est`, the `alarm` attribute has no effect. If the `alarm` timeout is reached, the hook terminates, but `pbs_est` continues to run.
- The `qrerun` command does not requeue job arrays. Instead, it leaves them in the *B* (begun) state.
- The `pbs.server.scheduler_restart_cycle()` function stops the current scheduler cycle, rather than restarting it. This behavior will change in a later release. When a hook's `fail_action` attribute is set to `scheduler_restart_cycle`, it uses the function, so the name of the action doesn't match the action.
- Do not perform partial removal of custom resources. It is important to remove custom resources correctly. See the instructions in the Resources chapter of the PBS Professional Administrator's Guide, section 5.14.2.
- Do not use different versions of PBS within the same complex. All machines using the same PBS server should run the same version of PBS. Do not mix major or minor versions. For example, do not run 13.1 and 14.2.3 in the same complex. While rolling upgrades are not supported, we do give some information on how to accomplish them; see [section 2.1.4, "Rolling upgrades", on page 6](#).
- Do not use `qmgr` to set or change the `vnode` sharing attribute, or a `vnode`'s `ncpus`, `vmem` or `mem`. Instead, create a new `vnode` definition file and use `"pbs_mom -s insert"` to add the file. Then HUP the MoM. There are require-

ments for the vnode definition file. See the PBS Professional Administrator's Guide, section 3.5, "How to Configure MoMs and Vnodes".

- All printable characters are allowed in a resource value, except for comma, space, colon, quotes, and ampersand, but are subject to some restrictions such as those imposed by `select` and `place`. Also, if these resource values are used as pathnames, they are subject to the restrictions of the operating system.
- Enabling X11 forwarding:

In order to enable X11 forwarding, add the path to the `xauth` binary to each MoM's `pbs_environment` file. For example:

```
echo $PATH
```

```
/bin:/user/bin
```

```
-bash-3.2$ which xauth
```

```
/usr/bin/X11/xauth
```

The entry in the `pbs_environment` file should be the following:

```
PATH=/bin:/usr/bin:/usr/bin/X11
```

- Make sure that all machines use the same Locale setting. If you try to modify the `Account` or `Project` attributes of an existing job using the `qalter` command or the `pbs_alterjob()` API so that the attribute will contain a non-ASCII, non-UTF-8 character, the PBS server and database may shut down. This happens when some machines are on a different Locale setting from others.
- If the `node_fail_requeue` server attribute is unset, it behaves as if it is set to zero, until the server is restarted. After the server is restarted, it reverts to its default value of 310.
- The functionality of `qdel -W force` has changed, and may cause unexpected behavior on some systems. See the `qdel` command in the PBS Professional Reference Guide.
- As of 13.0, PBS uses TCP instead of UDP for communication. See the Communication chapter in the Installation & Upgrading Guide.
- In the past, some PBS components defaulted to `/tmp` for storing temporary files. All components now default to `/var/tmp` which is most likely a persistent storage location. Administrators should take this into account and adjust the cleaning of `/var/tmp` accordingly.
- The exit value of an interactive job is always logged as zero in the accounting record, regardless of actual exit code value.
- When creating a MoM version 2 configuration file that will change natural vnode attributes, always explicitly specify `resources_available.ncpus`. Otherwise, vnodes accept only one job regardless of available CPUs.
- Warnings about using `pbsnodes` command:
 - a. You must define the `nmics`, `ngpus`, and `OS` resources in order to have their values displayed in the summary output.
 - b. While there is a header in the Job Summary output for suspended jobs on the vnode, this functionality is not yet implemented, and the value displayed is always zero.
- If you install a MoM but initially set `PBS_START_MOM=0` in `pbs.conf`, and start PBS, the MoM directories are not created. If you subsequently decide to run a MoM on this vnode, you must run `/opt/pbs/default/etc/pbs_habitat` in order to create the MoM directories.
- PBS-13383 (SPID 334661): If the `holidays` file is incomplete, `pbs_sched` hangs when `strict_ordering` and `backfill` are enabled.
- When using non-standard characters in file staging paths, make sure that your paths work with your shell and `scp` or `rcp`.

2.7 External Interface Changes

These are the external interface changes introduced in version 12.0 and newer.

2.7.1 Platform Deprecations

Table 2-1: Platform Deprecations

PBS Version	Change
12.0	Support for LAM MPI 6.5.9 is deprecated
12.0	Use of Python 2.5 is deprecated, to be replaced by 2.7 in a future release. (You can still use Python 2.5 with PBS 13.1.)
12.2	Support for HPCBP jobs is removed
14.2.1	Support for MPICH is deprecated
14.2.1	Support for MVAPICH is deprecated

2.7.2 Changes in Error Messages, Environment Variables, Etc.

Table 2-2: Changes in Error Messages, Environment Variables, Etc.

PBS Version	Change
12.0	Error message change: "Resources request is impossible to solve, job will never run" is replaced by "Error in calculation of start time of top job"
12.1	In <code>pbs.conf</code> , <code>PBS_CORE_LIMIT</code> is set to <i>unlimited</i> by default
13.0	Job and reservation names can be 236 characters
13.0	Fairshare and job sort formulas use Python math functions
13.0	The <code>TMPDIR</code> environment variable is deprecated and replaced with <code>PBS_TMPDIR</code>
14.2.1	New <code>pbs.conf</code> parameter <code>PBS_MOM_NODE_NAME</code>

2.7.3 Changes to Server and Scheduler Configuration Parameters

Table 2-3: Changes to Server and Scheduler Configuration Parameters

PBS Version	Change
12.2	The <code>smp_cluster_dist</code> scheduler parameter is deprecated
12.2	The <code>sort_queues</code> scheduler parameter is deprecated (this behavior is now automatic)
13.0	The <code>half_life</code> scheduler parameter is deprecated, replaced with <code>fairshare_decay_time</code> and <code>fairshare_decay_factor</code>
13.0	The <code>sort_queues</code> scheduler parameter has no effect.
13.0	The <code>resourcedef</code> file is deprecated (13.0)
13.1	The <code>preempt_priority</code> argument to the <code>job_sort_key</code> scheduler parameter is obsolete (this behavior is the default)
13.1	The <code>backfill</code> scheduler parameter is deprecated. Use <code>backfill_depth =0</code> instead.

Table 2-3: Changes to Server and Scheduler Configuration Parameters

PBS Version	Change
14.2.3	The <code>mom_resources</code> line in <code>sched_priv/sched_config</code> is deprecated. Instead, use <code>execheapost_periodic</code> hooks. See Example9-15: “Periodically update resources on vnodes” in the PBS Professional Hooks Guide.

2.7.4 Server and Scheduler Attribute Deprecations

Table 2-4: Server and Scheduler Attribute Deprecations

PBS Version	Change
	The server attribute <code>pbs_license_file_location</code> is deprecated, and replaced by <code>pbs_license_info</code> .
14.2.1	The <code>est_start_time_freq</code> server attribute is obsolete.
14.2.3	The <code>help_starving_jobs</code> scheduler attribute is deprecated. Instead, use <code>eligible_time</code> in the job formula. See also the <code>topjob_ineligible</code> job attribute.

2.7.5 Command Changes and Deprecations

Table 2-5: Command Changes and Deprecations

PBS Version	Change
	The <code>pbs_hostid</code> command is no longer available.
	The <code>-a</code> option to the <code>pbs_sched</code> command is deprecated.
12.0	New <code>pbs_interactive</code> command
12.1	The <code>pbs-report</code> command is deprecated and moved to the "unsupported" directory
12.1	<code>qmgr</code> command: question mark now equivalent to "help"
12.2	The <code>qsub -w interactive = true</code> option is deprecated; use <code>qsub -I</code>
12.2	You can edit the <code>qmgr</code> command line
13.0	You can create custom resources with <code>qmgr</code>
13.0	Added <code>qsub -f</code> option
13.0	The <code>xpbs</code> and <code>xpbsmon</code> interfaces to PBS are deprecated
13.0	Using <code>pbsrun_wrap</code> and <code>pbsrun_unwrap</code> for Intel MPI is deprecated
13.1	The <code>qorder</code> command is deprecated
14.2.1	The <code>pbs_probe</code> command is deprecated
14.2.1	New <code>qstat -E</code> option

2.7.6 New Server Attributes

Table 2-6: New Server Attributes

PBS Version	Change
13.0	<code>queued_jobs_threshold</code>
13.0	<code>queued_jobs_threshold_res.<resource></code>
13.1	<code>jobscript_max_size</code>
14.2.1	<code>python_restart_max_hooks</code>
14.2.1	<code>python_restart_max_objects</code>

Table 2-6: New Server Attributes

PBS Version	Change
14.2.1	python_restart_min_interval

2.7.7 Vnode Attribute Changes

Table 2-7: Vnode Attribute Changes

PBS Version	Change
11.3	New license_info attribute
12.2	The pcpus vnode attribute is deprecated

2.7.8 New Queue Attributes

Table 2-8: New Queue Attributes

PBS Version	Change
13.0	queued_jobs_threshold
13.0	queued_jobs_threshold_res.RES
13.1	backfill_depth

2.7.9 New Job Attributes

Table 2-9: New Job Attributes

PBS Version	Change
12.0	min_walltime
12.0	max_walltime
13.1	topjob_ineligible

2.7.10 New MoM Configuration Variables

Table 2-10: New MoM Configuration Variables

PBS Version	Change
12.0	\$reject_root_scripts

2.7.11 New Resources

Table 2-11: New Resources

PBS Version	Change
13.0	preempt_targets

2.7.12 New Scheduler Attributes

Table 2-12: New Scheduler Attributes

PBS Version	Change
13.0	sched_preempt_enforce_resumption
13.1	opt_backfill_fuzzy

2.7.13 New Hook Attributes

Table 2-13: New Hook Attributes

PBS Version	Change
13.0	fail_action
13.0	debug

2.7.14 New Hook Events

Table 2-14: New Hook Events

PBS Version	Change
13.0	execjob_launch
13.0	execjob_attach
13.0	exehost_startup

2.8 Documentation

Table 2-15: Documentation Changes

PBS Version	Change
14.2.1	New Hooks Guide; the Hooks Chapter in the Administrator's Guide has been removed.
13.1	Administrator's Guide contains a new Accounting chapter, and the previous Accounting chapter in the Reference Guide has been removed.
13.0	Installation and Upgrade Guide contains a new Communication chapter that describes the new pbs_comm daemon.
13.0	The Hooks chapter in the 13.0 Administrator's Guide has been overhauled. It also includes three new hook events.
12.1	User's Guide has been overhauled, with rewritten chapters on setting up the user's environment, submitting jobs, requesting resources, controlling the job, job arrays, and job file I/O.
12.0	Administrator's Guide chapter on Hooks is expanded and rewritten for the new execution-event hook feature. The "Setting File Transfer Mechanism" in the Administrator's Guide chapter on Administration is new for 12.0. Added "Initial Configuration" chapter to Installation & Upgrade guide for 12.0.

2.9 Hooks

- (13.0) PBS includes three new hook events and several other changes to the PBS Plugins infrastructure. See the Hooks chapter in the Administrator's Guide.
- Before creating any execution event hooks, read the Hooks chapter in the PBS Professional Administrator's Guide.
- (12.0) PBS allows hooks that run once a job is at the execution host, and hooks that periodically run on all execution hosts.
- Administrators who are writing provisioning hooks need to be aware that the Python `pbs.event().reject()` function fails to log the error message when used from a provisioning hook. This function behaves correctly for all other hooks supported by PBS Professional.
- Do not use the `pwd` module included with Python 2.5.1. This can cause a server crash.
- Support for Python 2.5 is deprecated in PBS 12.0. PBS will use a later version of Python in later releases of PBS.
- If a job tries to run where the complete update has not taken place, a hook could attempt to use resources that are not yet available, causing an unhandled exception. Make sure that your execution event hooks do not delete jobs if this

should happen. Write your hooks so that they trap unhandled exceptions and deal gracefully with the job. For example, you can use `pbs.event().job.rerun()`.

- As PBS copies or deletes execution or periodic hooks to the MoMs, the following messages are printed in the server's log file at 2047:

```
"successfully sent hook file <filename> to <mom_hostname>"
"successfully sent rescdef file <filename> to <mom host name>"
"successfully deleted hook file <filename> from <mom host name>"
"successfully deleted rescdef file <filename> from <mom host name>"
"failed to copy hook file <filename> to <mom host name>"
"failed to copy rescdef file <filename> to <mom host name>"
"failed to delete hook file <filename> from <mom host name>"
"failed to delete rescdef file <filename> from <mom host name>"
```

- The following are hook-related error codes:

Table 2-16: Hook-related Error Codes

Error Name	Code	Description
PBSE_MOM_INCOMPLETE_HOOK	15167	Execution hooks not fully transferred to a particular MoM
PBSE_MOM_REJECT_ROOT_SCRIPTS	15168	A MoM has rejected a request to copy a hook-related file, or a job script to be executed by root
PBSE_HOOK_REJECT	15169	A MoM received a reject result from a mom hook
PBSE_HOOK_REJECT_RERUNJOB	15170	Hook rejection requiring a job to be rerun
PBSE_HOOK_REJECT_DELETEJOB	15171	Hook rejection requiring a job to be deleted

- You may find that when you remove a hook, it takes some time for the hook to be completely purged. If you run `qmgr -c 'create hook <hook_name>'` where a previous hook of the same `<hook_name>` still exists, you will see the following message:

```
"hook name <hook_name> is pending delete, try another name"
```

 Either specify another name for the hook, or retry the `qmgr` request again later, after the previous hook is completely purged.
- You cannot create a hook configuration file using json format under Solaris or AIX.

2.10 SGI

- The SGI memacctd daemon is required for accurate accounting on shared-memory systems such as the UV. Without memacctd, PBS falls back to normal memory accounting.
- The `sgigenvodelist.awk` script is broken on any SGI systems using the cpuset MoM and running RHEL6, e.g. ICE and UV. If you don't make the fix, PBS will think that it can use cpus and mems that are already included in other cpusets, and MoMs will fail to come up. The error message indicating this problem looks like the following:


```
01/23/2017 13:29:14;0001;pbs_mom;Node;cpusets_initialize;/PBSPro (re)creation failed, CPU bits:
0000003f,ffffffff,ffffffff,ffffffff,ffffffff,ffffffff,ffffffff,fffffffe, mem bits:
00000000,00000000,00000000,00000000,00000000,00000000,00000000,00000000,00000000,00000000,00000000,00
000000,00000000,00000000,00000000,00000000,00000fff.
```

 Workaround: as root user, remove the trailing “-print” from line 118 in `$PBS_EXEC/lib/init.d/sgigenvodelist.awk`.
- Cpusets on SGI Systems:

RHEL 6.4 (and some 3.x-based Linux kernels) has changed the cpuset pseudofilesystem interface, by causing a prefix to be written to the cpuset controller's files with "cpuset" which messes up the cpuset code in PBS Professional. For example, the cpus file is now referred to as `/dev/cpuset/cpuset.cpus` (instead of `/dev/cpuset/cpus`). There is a "noprefix" flag to allow you to avoid breaking backwards compatibility by mounting just the cpuset controller on a hierarchy in the old way; it is silently used if you mount `-t cpuset`. But if you mount any unrelated cgroup controllers anywhere else, and especially if you also mount the cpuset controller in cgroups, the "noprefix" flag is ignored, even when you mount `-t cpuset`. One not-so-obvious way to cause this to happen is to enable RedHat's `cgconfig` tool which will premount cgroup hierarchies. Once `cgconfig` is running, then even if you use `"mount -t cpuset -o noprefix none /dev/cpuset"`, you'll get the prefixes. Bottom line, if you want to use the cpuset MoM on an SGI system, make sure that `/dev/cpuset` has filenames that are **not** prepended with 'cpuset.' . If they must be, make sure you don't run RedHat's `cgconfig` service and that you don't mount other cgroup controllers.
- Cpusets: the behavior of `pbs_mom` has changed with respect to actions taken when a CPU set setup failure occurs. See the documentation on the `cpuset_error_action` MoM configuration parameter.
- UV: PBS does not automatically generate vnode placement set information.
- When using SGI UV machines, make sure that the MTU size is consistent on all systems in the PBS complex.

2.11 MPI

- The default for starting wrapped mpids is `rsh`. If you are wrapping HP MPI with the PBS `pbs_mpihp`, and you want to use something other than `rsh`, be sure to set `PBS_RSHCOMMAND` in `PBS_HOME/pbs_environment`. For example:

```
PBS_RSHCOMMAND=ssh
```

- MPD rings:

As has been observed with `mpich2-1.0.3`, when starting MPD daemons, only one MPD per user on an execution host can be run; multiple MPD rings can be created per user as long as each participating execution host is in only one ring. Otherwise, jobs with intersecting execution hosts on their MPD rings can interfere with one another.

Suggested workarounds:

Given that only 1 MPD per user on an execution host can be run, do one of the following:

- Set 1 job per user to run on a PBS node by setting:

```
# qmgr -c "set node <node_name> max_user_run=1"
```

- Force users to submit jobs to run on nodes exclusively using the `excl` keyword in `-l place` as in:

```
% qsub -l select=2:ncpus=1:mpiprocs=2 -l place=free:excl job.scr
```

In addition, to set this default job-wide resource on the server, the admin can set:

```
# qmgr -c "set server resources_default.place=free:excl"
```

Note: Workaround 2 prevents other users from using the same node if it has more than 1 `resources_available.ncpus`.

- If you will use suspend/resume with Platform MPI, make sure that `MPI_BUNDLE_MPIDS` is not set.

2.12 Linux

2.12.1 Dependencies on Linux Shared Object Libraries

Altair distributes separate PBS packages for different supported Linux distributions. Customers running on SLES and RHEL distributions should take care to use the correct package for their systems. For example, in a complex made up of RHEL 6 and RHEL 7 hosts, the RHEL6 package should be used on the RHEL 6 hosts, and the RHEL7 packages should be used on the RHEL 7 hosts.

If PBS Professional 13.0 is deployed onto Linux distributions other than the supported SLES and RHEL distributions, there may be runtime problems with dependencies on shared object libraries, typically `libssl.so` or `libcrypto.so`.

Customers who encounter such problems should first determine which supported distribution is most similar to that installed on their systems. For example, `openSUSE 10.x` is similar to SLES 10, and CentOS distributions are similar to RHEL. The customer should then try the PBS package that is supported for the similar system.

If the runtime dependencies are still not resolved, there are two other approaches that may be tried. Before trying either approach, first run `ldd` on `pbs_mom` to get name and version number of libraries that PBS requires. `pbs_mom` needs to be extracted from the PBS rpm by using the Linux `rpm2cpio` and `cpio` commands.

Here is an example using an installed SLES 10 x86_64 pbs_mom:

```
shell> ldd /opt/pbs/default/sbin/pbs_mom
libdl.so.2 => /lib64/libdl.so.2 (0x00002b23b0a59000)
libcrypt.so.1 => /lib64/libcrypt.so.1 (0x00002b23b0b5d000)
libssl.so.0.9.8 => /usr/lib64/libssl.so.0.9.8 (0x00002b23b0c96000)
libc.so.6 => /lib64/libc.so.6 (0x00002b23b0ddd000)
libcrypto.so.0.9.8 => /usr/lib64/libcrypto.so.0.9.8 (0x00002b23b100d000)
/lib64/ld-linux-x86-64.so.2 (0x00002b23b093c000)
```

The Linux incompatibilities are generally associated with `libcrypto.so`, which is required by `libssl.so`.

Having identified the requisite version of `libcrypto.so` and `libssl.so`, there are two options:

1. Locate and install an OpenSSL compatibility library that will satisfy the dependency. In the example above, the library needs to satisfy the dependency on `libcrypto.so.0.9.8`.
2. Download and install a version of OpenSSL that corresponds to the dependency determined above (e.g., `libcrypto.so.0.9.8`).

3 Bugs

3.1 Modifications and Bug Fixes in PBS 14.2.3

Table 3-1: Bugs Fixed in PBS 14.2.3

JIRA ID	Summary
PBS-16681	pbs_est should time out after a configurable limit
PBS-17290	Sending large job scripts over TPP causes issues with TPP. Should use direct TCP connection in those cases.
PBS-17324	Tracejob output truncated on jobs using many vnodes
PBS-17535	mom's chk_file_sec messages missing when syslog is being used
PBS-17558	Upgrade ALSDK to fix server deadlock
PBS-17562	pbs.PERIODIC is documented but not defined in python module
PBS-17860	mom execjob_begin hook alarms when running a large multi-node job

3.2 Modifications and Bug Fixes in PBS 14.2.2

The following bugs are fixed in PBS 14.2.2:

Table 3-2: Bugs Fixed in PBS 14.2.2

JIRA Issue	Summary
PBS-15483	scheduler infinite loop after subjob fails to start on node
PBS-17802	Install test failure
PBS-17777	-fPIC error while linking with libpbs.a
PBS-17735	%preun commands run during an upgrade deregistering PBS as a service
PBS-17391	pbs_server has crashed (SIGABRT) when starting pbs_mom on a couple of racks (144 nodes) at the same time
PBS-17289	In high availability setup, "new" server leaf connection to pbs_comm keeps getting rejected because stale connection from old server is still seen
PBS-17280	Systemd PBS service startup times out if there are tens of thousands of jobs present
PBS-17205	logging via syslog does not work
PBS-17172	single or double quote in job env value, and the quote is not the last character in the value, causes qsub -V to lose env variable value as well as other variables when execjob_launch hook is present
PBS-17125	ACL and Array String attributes use memcpy for overlapping memory blocks
PBS-17105	only last SERVER hook of same kind (event) can modify PBS objects

Table 3-2: Bugs Fixed in PBS 14.2.2

JIRA Issue	Summary
PBS-15483	scheduler infinite loop after subjob fails to start on node
PBS-17074	qdel not working with seq_num.server
PBS-17047	sgigenvodelist.awk script broken on RHEL6
PBS-17005	jobs not getting backfilled when server backfill depth is zero but queue is non-zero
PBS-16562	Server unresponsiveness can hang scheduler while resuming a job during a qdel -Wforce

3.3 Known Bugs, Errors, and Warnings

This section lists known deficiencies and warnings as well as unresolved problem reports filed by customers.

Table 3-3: Open Bugs in PBS 14.2.3

JIRA Issue	Summary
PBS-7503	cp behavior differs from rcp/scp when sandbox=PRIVATE Workaround: When using sandbox=PRIVATE, use an absolute path in the "remote_path" portion, where the format is local_path@hostname:remote_path
PBS-7578	log_err with -l on Windows produces false success log message
PBS-9410	hook script cannot increase cput, walltime limits of running jobs even if qalter was executed by regular user
PBS-10149	printjob fails with 'malloc error' when used in prologue script.
PBS-10356	scheduler patiently waits forever if server_dyn_res script does not return on Linux Workaround: If a server_dyn_res script has a possibility of hanging, admins should take extra care to guard their scripts with alarms.
PBS-10699	Directories that are staged in are not being cleaned up after the job completes execution.
PBS-11046	Estimated start time stays set when job is held
PBS-11048	'Increasing running job walltime does not check for calendar conflicts
PBS-11156	epilogue timeout not getting logged with correct exit code
PBS-11177	Job array subjobs are not being held after failing/requeueing 20 times.
PBS-11220	Scheduler does not always resolve indirect resources
PBS-11338	qsig appears to signal more than just a session leader
PBS-11393	indirect resources will not resolve that point outside their host on complexes with multi-vnoded systems
PBS-11410	Attributes w/ unknown values in hook should be set to "None", not ""
PBS-11936	Add stronger distinction between pbs_license_info for socket vs LM-X
PBS-11940	resource named 'attributes' causes error when defined and used for a job

Table 3-3: Open Bugs in PBS 14.2.3

JIRA Issue	Summary
PBS-11963	PBS_NODEFILE on Cray sometimes shows compute nodes
PBS-12374	PBS cannot resolve job state inconsistency between server and MOM after node_fail_requeue
PBS-12387	if vnodedef sets the 'queue' node attribute, the associated queue's 'hasnodes' attribute will not be set. Workaround: Only set queue attribute via qmgr OR Set queue attribute via vnodedef file and then set reset one vnodes' queue attribute via qmgr
PBS-12456	pbs_mom should refuse connection when it comes from a different and non-failover pbs_server
PBS-12501	Nodes not marked offline for errors in new_cpuset that are not cpuset-lib API related
PBS-12531	If custom resource is named the same as formula keyword, formula will quietly use keyword
PBS-12683	Interactive job remains in R state forever after failed attempt to connect to job's host on windows
PBS-12689	When resuming a suspended job, the scheduler will recalculate but not alter the walltime
PBS-12769	incorrect node_sort_key entry causes scheduler to core dump (and not start up)
PBS-12813	The hook on execjob_end times out when using "pbs.server()"
PBS-12821	PBS Pro writing garbage characters in to accounting log entry which makes that record totally unusable for Analytics
PBS-12880	pbsnodes -l slowness
PBS-12950	CPUSet not deleted if there are network issues during end of job processing Recommendation: Once identified, stop MOM with SIGINT, remove zombie from mom_priv/jobs, restart UV MOM with -p flag
PBS-12996	MOM-server/scheduler sync issue following overlay upgrade re: which jobs are "done" and can have cpuset destroyed.
PBS-13001	MOM in boot cpuset on linux_cpuset machines breaks socket licensing as "whole system" flag is not passed to hwloc Workaround: Temporarily start MOM in the root cpuset, and move it to the boot cpuset later.
PBS-13028	PBS consumes floating licenses for jobs which are running on socket licensed vnodes after floating license file update
PBS-13081	When using v2 config to change "natural node" attribute node becomes job-busy even when only 1 cpu is used. Workaround: Always explicitly specify resources_available.ncpus in a config v2 file that is meant to change natural node attributes.
PBS-13125	potential scheduler crash in reorder_nodes()
PBS-13174	qsub -Wblock=true hangs server for 190 seconds per job at exit if submission host is unresponsive

Table 3-3: Open Bugs in PBS 14.2.3

JIRA Issue	Summary
PBS-13180	Dataservice won't create if you have "=" in PBS directory name
PBS-13247	susp field is not getting updated in pbsnodes -SjL
PBS-13249	Mom hook runs periodically after the event is removed in qmgr
PBS-13260	the command <code>pbs.server().scheduler_restart_cycle()</code> does not restart the scheduling cycle. Workaround: Wait for a job to enter or exit the system.
PBS-13383	<code>pbs_est</code> hangs if holiday file is not complete Workaround: Add missing days.
PBS-13384	<code>pbs_sched</code> hangs with <code>strict_ordering</code> and <code>backfill</code> enabled - holidays file incomplete Workaround: Add missing days.
PBS-13406	hook throws exception if it tries to access <code>vnodes</code> and none are defined
PBS-13448	PBS assigns jobs to offlined MoMs
PBS-13476	Reinstalling PBS assumes PBS Datastore user account to be <code>pbsdata</code> Workaround: <code>rm -rf /etc.pbs.conf /etc/init.d/pbs \$PBS_HOME \$PBS_EXEC</code>
PBS-13490	Lack of <code>PBS_CORE_LIMIT</code> in <code>/etc/pbs.conf</code> sets <code>coredumpsize</code> of job processes (soft and hard) to 0 Workaround: Specify <code>PBS_CORE_LIMIT</code> in <code>/etc/pbs.conf</code>
PBS-13517	IntelMPI wrapper support job launch fails when submitting large jobs
PBS-13539	full <code>exec_host</code> missing in <code>qstat -n</code> output when <code>vnode</code> name contains period character
PBS-13584	<code>tm_attach</code> (<code>pbs_attach</code>) is unable to get the process user id when forked process is still owned by root
PBS-13603	Epilogue running at the same time as <code>execjob_end</code> for the same job
PBS-13639	<code>pbs_mom</code> is too quick to abort a job
PBS-13660	Can't access reservation <code>Duration</code> attribute in hook
PBS-13706	Preempted jobs fails to resume when the high priority job finishes because another job was scheduled on its resources
PBS-13712	unable to submit jobs on *nix if stagein filename contains '@' symbol
PBS-13715	Zombie jobs from reimaged nodes (when MOM loses <code>PBS_HOME</code> or gets corrupted files in jobs directory) Recommendation: Discover a zombie because users notice that the other jobs don't run, and either <code>qrerun</code> it (for rerunnable jobs) or <code>qdel</code> it.

Table 3-3: Open Bugs in PBS 14.2.3

JIRA Issue	Summary
PBS-13725	cpuset MOM fails to start exclhost multihost jobs Workaround: Don't use -lplace=exclhost when targeting cpuset MOM clusters
PBS-13752	qalter, qmove hooks get an internal PBS error with jobs using "-X".
PBS-13792	PBS should better handle cpuset MOM resource availability that persistently does not match actual cpuset resource availability resulting in cpuset creation failures
PBS-13798	PBS has changed the contents of \$PBS_NODEFILE, now lists FQDN
PBS-13809	PBS Server hang, seemingly in internal license accounting
PBS-13843	Address PP-60: Using kill -HUP 'pgrep pbs_mom' could result in signal being sent to children of main MOM
PBS-13873	Job array subjob state and array_indices_remaining can become inconsistent Workaround: Change the state in the datastore to make the subjob state and array_indices_remaining attribute consistent
PBS-13887	Running jobs lose stime attribute after server restart
PBS-13946	db upgrade fails when db user is different from original db creator Workaround: If you have a backup of PBS_HOME: Delete the new busted PBS_HOME. Copy in the old one, point the \$PBS_EXEC default link back at the old version, look in the old datastore to see what the old account was (see note in request 359102 for details), re-create the user account, run pbs_ds_password -r -C pbs (the old user name), then remove the 13.0 rpm and install again.
PBS-13976	Resources are reserved for starving jobs ahead of reserving them for resuming the preempted job
PBS-13992	Jobs in a standing reservation are starting later than expected after the begin time of the reservation was changed Workaround: Delete and re-create the standing reservation
PBS-14013	PBS_CONF_REMOTE_VIEWER configuration parameter (supported only on Windows) is not working appropriately
PBS-14040	corrupt MoM hook python file not updated until MoM restarted
PBS-14051	High priority job which needs to preempt other jobs relevant to a limit cannot run through preemption if preemption candidates are using different resources.
PBS-14059	Server and sched treat jobs in exiting state as "running" wrt resources_assigned vs. resources_available, but sched treats those jobs as "not running" wrt limits checks such as max_run_res. Workaround: Use resources_available for limit checking, but only works for PBS_ALL.
PBS-14111	PBS does not check for limits.pbs_comm file Workaround: Add a few lines to /etc/init.d/pbs

Table 3-3: Open Bugs in PBS 14.2.3

JIRA Issue	Summary
PBS-14223	Description for node_sort_key in sched_config is incorrect
PBS-14225	Hook error with queue.resources_available.keys()
PBS-14294	unable to submit jobs on Windows if stagein filename contains '@' symbol
PBS-14300	Scheduler hung while calendaring
PBS-14401	Execution node get autocreated even though it was deleted through qmgr
PBS-14421	Email is not getting triggered when a job fails during stage-in/stage-out
PBS-14422	Limits set "for SGI" in limits.pbs_<daemon> shipped should apply to _all_ Linux machines Workaround: Fix the limits.pbs_<daemon> files.
PBS-14494	Primary server takes a long time to come back up, after failover, when lots of jobs are in the queue
PBS-14528	cannot mix queued_jobs_threshold on server and queues
PBS-14798	get_wm should not call dlsym repeatedly (i.e. for every process and every MOM poll cycle)
PBS-14820	server will almost always fail when attempting to abort a running job during start up
PBS-14833	Job script limit is lower than 1GB due to PostGreSQL, not 2GB Workaround: Avoid setting jobscript_max_size to something so large that it doesn't work.
PBS-14856	race condition in mom hook transport initiated by provisioning
PBS-14871	qstat -Qw gives wrong output
PBS-14882	PBS leaves Unix socket file lying around when TMPDIR too long.
PBS-14902	server crash in pbs_python_mark_object_readonly
PBS-14998	Custom resources created in mom hook can sometimes cause server crash when mom sends job obit
PBS-14999	PBS Fairshare: PBS fairshare stores share value as a negative number upon buffer overflow
PBS-15000	Small memory leak in mom_topology
PBS-15109	Reservations display poorly if -R -D, and -E are given to pbs_rsub Workaround: Only specify two of -R, -D, and -E
PBS-15117	MoM not waiting for ping time before getting marked down
PBS-15121	pbs_mom can hang timing out interactive job if job process froze writing to terminal
PBS-15123	scheduler infinite loop when trying to run job on "free" node with no mom running
PBS-15134	Slow but persistent memory leaks in scheduler in find_jobs_to_preempt Workaround: Kill and restart the scheduler (and issue qmgr -c "s scheduling = true") when the RSS grows over a certain size.
PBS-15277	slight memory leak at mom

Table 3-3: Open Bugs in PBS 14.2.3

JIRA Issue	Summary
PBS-15284	pbs_sched respawn on SEGV fails when abort() of original scheduler takes longer than 10 seconds Workaround: Run a cron job doing /etc/init.d/pbs staert (or service pbs start, if appropriate, depending on platform) regularly.
PBS-15311	Jobs become zombies when the node is reimaged, the mom loses PBS_HOME or there are corrupted files in the job directory
PBS-15402	PBSPro failover secondary server fails to continuously check whether it needs to start a scheduler locally
PBS-15423	umask is not overridden when used in default_qsub_arguments
PBS-15437	False failovers during job recovery caused by delays to do logging Workaround: Use softlinks to ensure that even when PBS_HOME is shared, the server log files are stored on local disk.
PBS-15484	node comments lost after pbs_server restart
PBS-15565	lock files for mom and sched exist (with wrong permissions) even when the daemons are not running
PBS-15568	job ignores the delete request and rerun by scheduler (race condition)
PBS-15570	Cannot update estimated_start/vnodes job attributes of a suspended job Workaround: Increase max_starve to a value high enough so that the preempted jobs resume before being marked as starving.
PBS-15572	Scheduler incorrectly logs that jobs in reservation queues do not exist any more
PBS-15575	Need to reduce server quasihang in vnode_available when checking 1000 instances of standing reservations
PBS-15577	deleted custom resources which are added by exechoost_startup hook reappearing again after sending SIGHUP signal to mom
PBS-15578	sub job not resuming after suspended via qsig
PBS-15579	Modifying multiple jobs with the same qalter command clobbers -W depend list Workaround: qalter jobs one at a time
PBS-15604	Failover after power loss on primary makes secondary try to connect to data service on primary -- never gives up.
PBS-15670	Wrong value for cput being updated at the end of job in qstat -fx output
PBS-15721	qsub accepts invalid arguments in PBSPro_14.2.0
PBS-15961	Non-rerunnable job is not getting aborted as expected (intermittent)
PBS-15963	pbsdsh system errors when scaling to 1024 connections
PBS-15965	Segfault in qsub when cwd removed
PBS-16044	Address PP-448: max_run_res limits cannot be read from within queuejob hook

Table 3-3: Open Bugs in PBS 14.2.3

JIRA Issue	Summary
PBS-16045	Address PP-449: hook debug .data file may contain multiple entries on a single line
PBS-16091	invisible resources are printed in the server_logs and tracejob
PBS-16139	Queue resources_assigned.ncpus becomes negative and remains even after unset the attribute.
PBS-16368	Long Reservation Length goes negative
PBS-16438	Scheduler loses one cycle worth of fairshare data when it receives SIGHUP
PBS-16461	Server crashes if mom hook sets resources_used on a non-existent resource
PBS-16486	execjob_launch hook is not getting rejected if job attribute is modified
PBS-16641	Message "Job Modified at request of Scheduler" not present in server log on simple job submission
PBS-16736	Secondary PBS server sometimes fails to stop
PBS-16783	log message about warning while reservations get confirmed is not visible anymore in server_logs
PBS-16863	PBS_MOM_NODE_NAME not resolving hostname on windows
PBS-16893	If PBS_START_* daemon variables are exported in installation environment they will be missing from pbs.conf on fresh install
PBS-16895	Windows:Primary and Secondary server take more time to startup than usual from a shared home
PBS-16896	Force rerunning a job twice results in the second 'R' record to not have resource_used information
PBS-16961	In HA setup, "new" server leaf connection to pbs_comm keeps being rejected because stale connection from old server is still seen
PBS-16962	<p>Routing queues do not guarantee in-order routing, cause races when dependencies exist</p> <p>Workaround:</p> <p>No easy one. Admin intervention to temporarily raise limits for the user is usually necessary to break the deadlock.</p>
PBS-16963	When large number of jobs are submitted.few of them remains in Q state with comment "Execution server rejected request"
PBS-17009	resource limits on pulling queue not working when using local peering
PBS-17028	node_fail_requeue causes duplicate accounting records for job
PBS-17030	Cray Bug 847878 - /etc/init.d/pbs start: line 424: echo: write error: No such process
PBS-17035	qdel not working with seq_num.server
PBS-17046	Scheduler may not mark a node for provisioning when it should be
PBS-17078	alarm does not stop server periodic hook as it does other hooks
PBS-17081	If the scheduler can't calculate the start time for a top job, it still takes a top job slot
PBS-17092	Job in substate 41 can't be deleted while in transit to MOM
PBS-17166	qmgr output doesn't have license info as it was in PBS 13.X
PBS-17170	When using qstat -w -T "Est" part of label in header does not line up correctly

Table 3-3: Open Bugs in PBS 14.2.3

JIRA Issue	Summary
PBS-17180	triple quotes in job environment cause job to fail to launch when there is an execjob_launch hook present
PBS-17211	qmgr is not case sensitive towards setting limits on custom resources but scheduler is case sensitive
PBS-17213	pbs_python unable to use hook input for debugging of runjob hook event
PBS-17220	Environment variable with comma is not passed properly to job using qsub -v
PBS-17318	Windows pbs_mom.exe crash in check_pwd()
PBS-17319	PBS 14.2.1 install overrides site-selected accounting log permissions
PBS-17332	qmgr allows wild cards for users in the acl list
PBS-17462	Upgrade from 13.1.0 to 14.2.1 failed with pbs_ds_password: PBS dataservice authentication failed
PBS-17506	static server/queue level resource is incorrectly accounted if jobs run within reservations
PBS-17522	number of queued jobs value is negative
PBS-17536	PBS does not cleanup after sloppy hooks
PBS-17579	misuse of qrun -H crashes server
PBS-17651	Missing "Successfully deleted hook file..." message from server logs
PBS-17736	memory corruption in strncpy (__strncpy_sse2_unaligned), possible server hang
PBS-17797	RPM upgrades from 14.2.1 are problematic Workaround: Avoid using “rpm -U”. Instead, uninstall 14.2.1 using “rpm -e”, then install the next version using “rpm -i”.
PBS-17811	PBS Interactive job shell does not start on a Windows multinode machine
PBS-17830	Unable to submit standing reservations if the actual time is between what was specified as -R and -E and gives error "pbs_rsub: Bad time specification(s)"
PBS-17844	PBS_NODEFILE doesn't include all vnodes
PBS-17859	Address PP-772: all mom hook events run in a single iteration of the main mom loop are counted against job's walltime
PBS-17896	Not all calls to pbs.server() in cgroup hooks are guarded with dotimeout(10,...)
PBS-17986	Unable to debug server periodic hooks with pbs_python
PBS-18041	Address PP-822: Client not accepting response due to slow network hangs server
PBS-18060	Hooks run on MoM before server has sent updated resourcedef file can crash in _svr_types.py (uncaught exception)
PBS-18078	Variable_List corruption when job is read from .JB file from disk
PBS-18098	Job accrue type is incorrect after dependency is released
PBS-18133	Negative total job count in qstat -Q
PBS-18136	Interactive jobs spamming their stdout/stderr can hang MoM when it tries to kill the job
PBS-18153	Jobs with multiple exit records in accounting log

Table 3-3: Open Bugs in PBS 14.2.3

JIRA Issue	Summary
PBS-18172	MIN_STACK_LIMIT for stacksize not working for PBS Pro version 14.x.x

Legal Notices

The following copyrights apply:

4.1 avltree

AVL index routines by Gregory Tseytin

Copyright (c) 2000 Gregory Tseytin <tseyting@acm.org>. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer as the first lines of this file unmodified.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY Gregory Tseytin ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED.

IN NO EVENT SHALL Gregory Tseytin BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

4.2 popen.c

Copyright (c) 1988, 1993 the Regents of the University of California. All rights reserved.

This code is derived from software written by Ken Arnold and published in UNIX Review, Vol. 6, No. 8.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgement:
This product includes software developed by the University of California, Berkeley and its contributors.
4. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

\$FreeBSD: src/lib/libc/gen/popen.c, v 1.14 2000/01/27 23:06:19

4.3 Cray

Copyright (c) 2003-2010 Cray Inc. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name Cray Inc. nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

4.4 libical

The contents of this file are subject to the Mozilla Public License Version 1.1 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at <http://www.mozilla.org/MPL/>

Software distributed under the License is distributed on an "AS IS" basis, WITHOUT WARRANTY OF ANY KIND, either express or implied. See the License for the specific language governing rights and limitations under the License.

The Original Code is libical.

The Initial Developer of the Original Code is Eric Busboom

All Rights Reserved.

Contributor(s): See individual source files.

4.5 Expat

Copyright (c) 1998, 1999, 2000 Thai Open Source Software Center Ltd and Clark Cooper

Copyright (c) 2001, 2002, 2003, 2004, 2005, 2006 Expat maintainers.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT.

IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

4.6 hwloc

Copyright © 2009 CNRS

Copyright © 2009 INRIA. All rights reserved.

Copyright © 2009 Université Bordeaux 1

Copyright © 2009 Cisco Systems, Inc. All rights reserved.

See COPYING in top-level directory.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. The name of the author may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE AUTHOR "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

4.7 editline

Libedit license (BSD style): <http://sourceforge.net/projects/libedit/>

Copyright (c) 1992, 1993 The Regents of the University of California. All rights reserved.

This code is derived from software contributed to Berkeley by Christos Zoulas of Cornell University.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

4.8 MinGwEditLine

MinGwEditLine: <http://sourceforge.net/projects/mingweditline/>

A BSD-licensed EditLine API implementation for the native Windows Console

Copyright © 2010-2012 Paolo Tosco

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. Neither the name of MinGWEditLine nor the name of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDER AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Index

`$reject_root_scripts` [16](#)

A

account

 data service [8](#)

Accounting chapter [17](#)

AIX [19](#)

almutil [8](#)

Altair License Server [2](#), [7](#)

B

backfill [12](#)

backfill_depth [12](#), [16](#)

C

CentOS [21](#)

 6.x and 7.x [1](#)

cluster products [1](#)

Communication chapter [17](#)

cpuset pseudofilesystem interface [20](#)

cpuset_error_action [20](#)

Cray

 upgrade [5](#)

 which PBS version to use [1](#)

CSA [2](#)

custom resources

 removing [9](#)

D

data service account [8](#)

database

 upgrade [6](#)

debug [17](#)

docs

 Accounting chapter [17](#)

 Communication chapter [17](#)

 Hooks chapter [17](#)

 User's Guide [17](#)

E

exechost_startup [17](#)

execjob_attach [17](#)

execjob_launch [17](#)

exit value of interactive job [10](#)

F

fail_action [9](#), [17](#)

FLEX licenses [7](#)

floating licenses [7](#)

H

half_life [12](#)

holidays file [10](#)

hook

 deleting [19](#)

 error codes [19](#)

 provisioning [18](#)

 updates [18](#)

Hooks chapter [17](#)

HP MPI [2](#), [21](#)

HPCBP [12](#)

I

ICE [1](#)

Intel MPI [2](#), [14](#)

J

job_sort_key [12](#)

jobscript_max_size [14](#)

json [19](#)

L

LAM MPI

 6.5.9 [12](#)

LAM/MPI [2](#)

libcrypto.so [21](#)

libdl [8](#)

libssl.so [21](#)

license_info [16](#)

licenses

 FLEX [7](#)

 floating [7](#)

 socket [7](#)

Locale [10](#)

M

max_walltime [16](#)

migration upgrade [5](#)

min_walltime [16](#)

MoM version 2 configuration file [10](#)

Index

MPD rings [21](#)
MPI_BUNDLE_MPIDS [21](#)
MPICH [1](#)
MPICH2 [2](#)
MPICH-GM [2](#)
MTU [20](#)
MUNGE [2](#), [8](#)
MVAPICH [2](#)
MVAPICH2 [2](#)

N

ngpus [10](#)
nmics [10](#)
node_fail_requeue [10](#)

O

OpenMPI [2](#)
opt_backfill_fuzzy [17](#)
OS resource [10](#)
overlay upgrade [5](#)

P

PBS 13.0.40x [1](#)
PBS Works Documentation Download Area [2](#)
pbs.server.scheduler_restart_cycle() [9](#)
PBS_CORE_LIMIT [12](#)
pbs_habitat [10](#)
pbs_hostid [14](#)
pbs_interactive [14](#)
pbs_license_file_location [7](#), [14](#)
pbs_license_info [7](#), [14](#)
pbs_license_linger_time [7](#)
pbs_license_min [7](#)
pbs_mom [20](#)
pbs_mpihp [21](#)
PBS_RSHCOMMAND [21](#)
pbs_sched [14](#)
PBS_TMPDIR [12](#)
pbsdata [8](#)
pbsnodes [10](#)
pbs-report [14](#)
pbsrun_unwrap [14](#)
pbsrun_wrap [14](#)
pcpus [16](#)
Platform MPI [2](#), [21](#)
preempt_priority [12](#)
preempt_targets [16](#)
provisioning hook [18](#)
pwd Python module [18](#)
Python
 2.5 [12](#)
 math functions in fairshare and formula [12](#)
 pwd [18](#)

Q

qdel -W force [10](#)
qmgr [14](#)
qorder [14](#)
qsub
 -f [14](#)
 -W interactive = true [14](#)
queued_jobs_threshold [14](#), [16](#)
queued_jobs_threshold_res [14](#), [16](#)

R

Red Hat Enterprise Linux [1](#)
removing custom resources [9](#)
resource
 ngpus [10](#)
 nmics [10](#)
 OS [10](#)
resourcedef [12](#)
resources
 values [10](#)
resources_available.ncpus [10](#)
RHEL [21](#)
rolling upgrade [6](#)

S

sched_preempt_enforce_resumption [17](#)
scheduler_restart_cycle [9](#)
SGI MPI (MPT) [2](#)
SGI Performance Suite [1](#)
sharing
 setting [9](#)
SLES [21](#)
smp_cluster_dist [12](#)
socket licenses [7](#)
Solaris [19](#)
sort_queues [12](#)
storing temporary files [10](#)
SuSE [1](#)

T

TMPDIR [12](#)
topjob_ineligible [16](#)

U

unsupported [14](#)
upgrade
 Cray [5](#)
 database [6](#)
 migration [5](#)
 overlay [5](#)
 rolling [6](#)
User's Guide [17](#)

Index

UV [1](#), [20](#)

V

versions of PBS in same complex [9](#)

X

X11 forwarding [10](#)

xpbs [14](#)

xpbsmon [14](#)

Index
