

# Table of Contents

<b>How to check which batch manager is used</b> .....	1
<b><i>Determine which batch manager is used</i></b> .....	1
<b><i>High level check of the status of each batch system</i></b> .....	3
SLURM .....	3
PBSPro .....	3
TORQUE .....	5
SGE .....	5
LSF .....	6



# How to check which batch manager is used

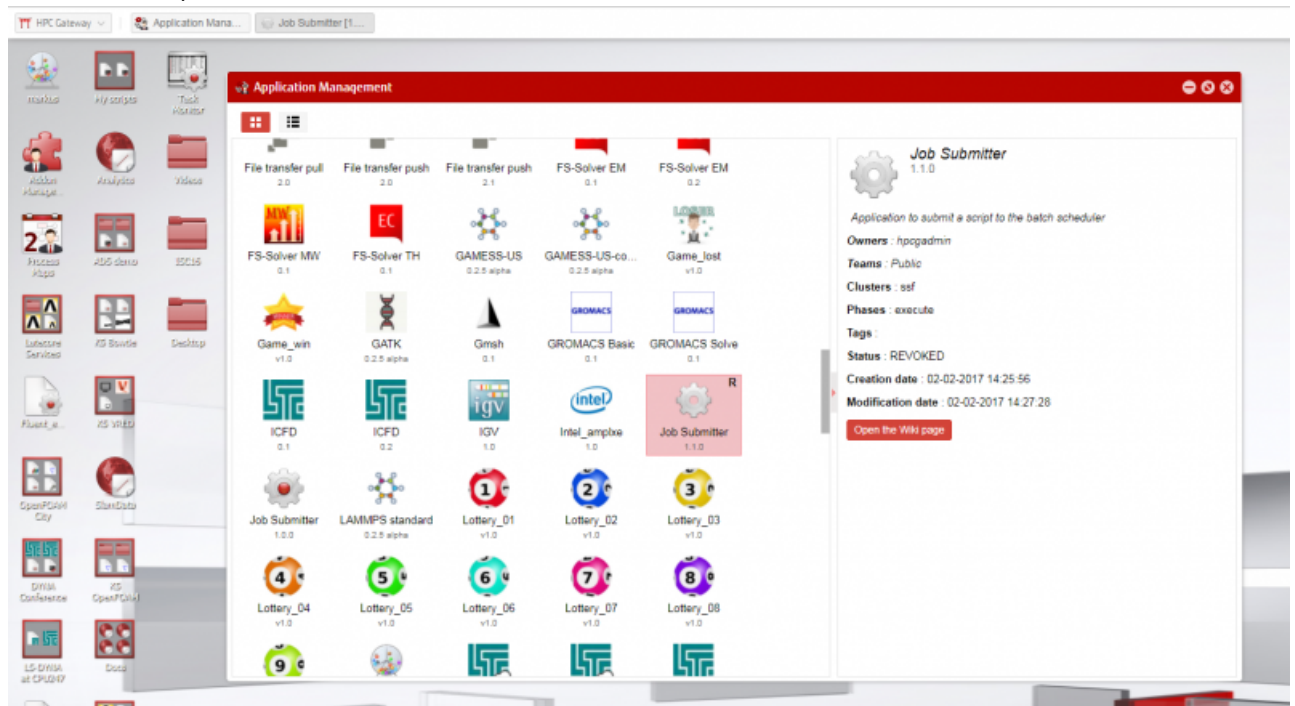
The Gateway supports many different batch managers. When job submission fails it is important to determine the batch manager is responding correctly.

On this page we describe how to do that for a variety of batch manager systems.

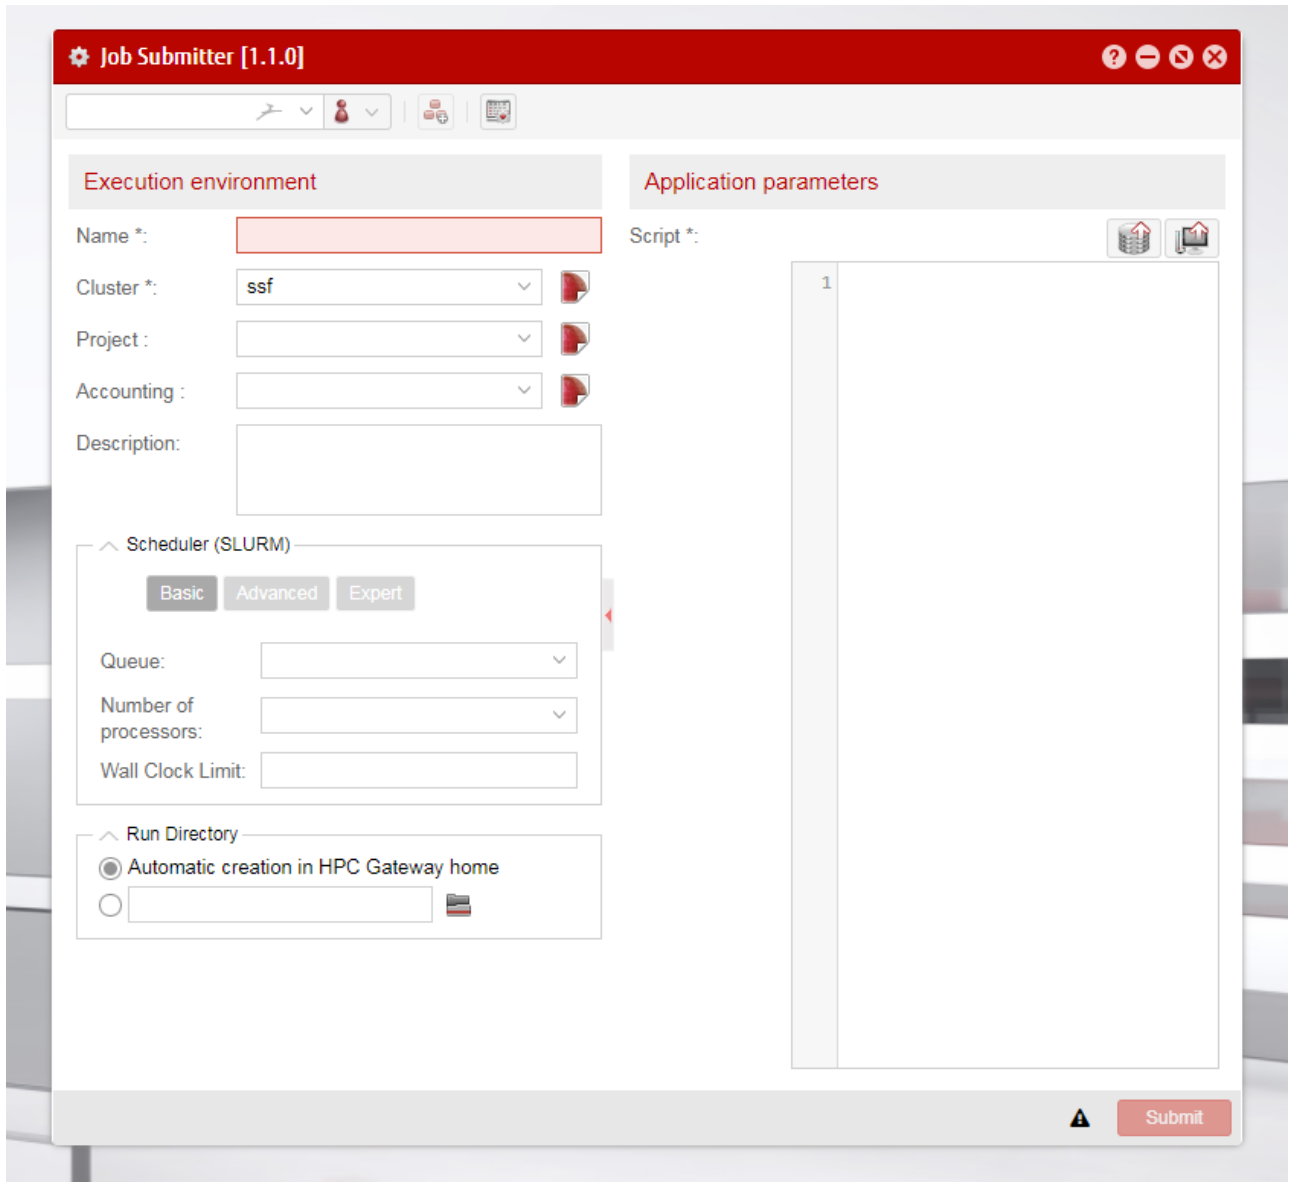
## Determine which batch manager is used

If the user does not know the batch manager type being used then the Gateway Portal can be used to check which batch system is used.

1. Ask the user to login to the Portal and choose the Application Manager from the “start” menu in the top left (the button with 'HPC Gateway' text on it)
2. Then ask them to scroll until they see the “Job Submitter” icon ( it uses the task wheel icon as shown below).



3. Get them to double click on this icon to start the “Job Submitter”
4. In the Job submitter windows get them to select the correct “Cluster” which is a parameter of the Execution environment parameters on the left of the window.



5. The batch manager name will be displayed next to the right of the “Scheduler” section about half way down on the left.  
In this example (above) we can see that the SLURM batch scheduler has been configured for the 'ssf' cluster.

Gateway supports the following batch managers:

- [SLURM](#)
- [PBS Professional \(or PBSPro\)](#)
- [Torque](#)
- [SGE \(Sun Grid Engine and/or its derivatives\)](#)
- [LSF \(Load Share Facility\)](#)

Confirming the status of each batch system is different. See the next section to run a very simple high level check for each batch manager type.

## High level check of the status of each batch system

### SLURM

Run the 'sinfo' command to check if the system is operating correctly.

```
[hpcgadmin@ssf root]$ sinfo
PARTITION AVAIL  TIMELIMIT  NODES  STATE NODELIST
all*       up 1-00:00:00      2 down* cn[05,07]
all*       up 1-00:00:00      2 drain cn[03-04]
all*       up 1-00:00:00     12  idle cn[01-02,06,08-16]
fujitsu    up 1-00:00:00      2 drain cn[03-04]
fujitsu    up 1-00:00:00      2  idle cn[01-02]
mbda       up 1-00:00:00      2 down* cn[05,07]
mbda       up 1-00:00:00      2  idle cn[06,08]
intel      up 1-00:00:00      8  idle cn[09-16]
gateway    up   1:00:00       2 down* cn[05,07]
gateway    up   1:00:00       2 drain cn[03-04]
gateway    up   1:00:00     12  idle cn[01-02,06,08-16]
[hpcgadmin@ssf root]$
```

The system should respond with configuration and status of the defined batch partitions as shown above.

If the command fails or does not produce similar output there is likely a problem.

### PBSPro

For the PBSPro subsystem you can usually determine whether the main PBS daemon is running by executing the "qstat" command.

```
[hpcgadmin@autan ~]$ qstat -a

autan.fujitsu.fr:

                               Req'd  Req'd
Elap                               S
Job ID      Username Queue   Jobname  SessID  NDS  TSK  Memory  Time  S
Time
-----
284.autan.fujit hpcgadmi workq   caffeframe  --     1   1    --    --   H
--
[hpcgadmin@autan ~]$
```

If the system is working correctly the command will return something similar to the above (though the list of jobs could be very long), or it will fail with a message "could not connect to server".

Additionally you can use the "pbsnodes -a" command to determine the status of all compute nodes of

the cluster. Simply run this command to see if the server is responding and the state of all nodes.

```
[hpcgadmin@autan ~]$ pbsnodes -a
knl04
  Mom = knl04.default
  ntype = PBS
  state = down
  pcpus = 272
  resources_available.arch = linux
  resources_available.host = knl04
  resources_available.mem = 98875016kb
  resources_available.ncpus = 272
  resources_available.queue_name = workq
  resources_available.vnode = knl04
  resources_assigned.accelerator_memory = 0kb
  resources_assigned.mem = 0kb
  resources_assigned.naccelerators = 0
  resources_assigned.ncpus = 0
  resources_assigned.vmem = 0kb
  comment = node down: communication closed
  resv_enable = True
  sharing = default_shared
  license = l

knl05
  Mom = knl05.default
  ntype = PBS
  state = free
  pcpus = 272
  resources_available.arch = linux
  resources_available.host = knl05
  resources_available.mem = 98875016kb
  resources_available.ncpus = 272
  resources_available.queue_name = workq
  resources_available.vnode = knl05
  resources_assigned.accelerator_memory = 0kb
  resources_assigned.mem = 0kb
  resources_assigned.naccelerators = 0
  resources_assigned.ncpus = 0
  resources_assigned.vmem = 0kb
  resv_enable = True
  sharing = default_shared
  license = l

.
.
.
.
```

The output can be quite long if the site has many compute nodes. The important thing is the

command responds and prints output similar to the above.

## TORQUE

For TORQUE you can use the same commands as for PBSPro as it was a fork of this code many years earlier. However the commands have been largely kept compatible.

So run the 'qstat -a' and the 'pbsnodes -a' as for PBSPro to get the status of a Torque cluster.

## SGE

The status of an SGE batch manage can be obtained by executing the the 'qhost' and 'qstat' commands.

### Host/Node Status: qhost

Node or host status can be obtained by using the qhostcommand. An example listing is shown below.

```
[hpcgadmin@autan ~]$ qhost
```

HOSTNAME	ARCH	NPROC	LOAD	MEMTOT	MEMUSE	SWAPTO	SWAPUS
global		-		-	-	-	-
node000		lx24-amd64	2	0.00	3.8G	35.8M	0.0
node001		lx24-amd64	2	0.00	3.8G	35.2M	0.0
node002		lx24-amd64	2	0.00	3.8G	35.7M	0.0
node003		lx24-amd64	2	0.00	3.8G	35.6M	0.0
node004		lx24-amd64	2	0.00	3.8G	35.7M	0.0

### Queue Status: qstat

Queue status for *jobs* can be found by issuing a qstat command. An example qstat issued by user *deadline* is shown below.

```
[hpcgadmin@autan ~]$ qstat
```

job-ID	prior	name	user	state	submit/start	at	queue	slots	ja-task-ID
304	0.60500	Sleeper4	deadline	r	01/18/2008	17:42:36			
cluster@norbert	4								
307	0.60500	Sleeper4	deadline	r	01/18/2008	17:42:37			

```
cluster@norbert 4
310 0.60500 Sleeper4 deadline qw 01/18/2008 17:42:29
4
313 0.60500 Sleeper4 deadline qw 01/18/2008 17:42:29
4
316 0.60500 Sleeper4 deadline qw 01/18/2008 17:42:29
4
321 0.60500 Sleeper4 deadline qw 01/18/2008 17:42:30
4
325 0.60500 Sleeper4 deadline qw 01/18/2008 17:42:30
4
308 0.53833 Sleeper2 deadline qw 01/18/2008 17:42:29
2
```

## LSF

Use the 'bhosts' command to see whether the LSF batch workload system is running properly.

### bhosts command

The bhosts command displays the status of LSF batch server hosts in the cluster, and other details about the batch hosts:

- Maximum number of job slots that are allowed by a single user
- Total number of jobs in the system, running jobs, jobs that are suspended by users, and jobs that are suspended by the system
- Total number of reserved job slots

Normal operation requires that hosts have the status ok .

```
% bhosts
HOST_NAME      STATUS      JL/U      MAX      NJOBS      RUN      SSUSP      USUSP
RSV
hosta          ok          -         -         0          0         0          0
0
hostb          ok          -         -         0          0         0          0
0
hostc          ok          -         -         0          0         0          0
0
hostd          ok          -         -         0          0         0          0
0
```

If you see the following message then LSF is not working correctly and the Gateway will not be able to submit jobs.

```
batch system daemon not responding ... still trying
```



Note: If you have just started or reconfigured LSF, wait a few seconds and try the 'bhosts' command again to give the 'mbatchd' daemon time to initialize.

From:

<https://docs.fujitsu.fr/HPCGateway/current/> - HPC Gateway

Permanent link:

[https://docs.fujitsu.fr/HPCGateway/current/doku.php?id=fujitsu:hpcgateway:support:check\\_batch\\_manager](https://docs.fujitsu.fr/HPCGateway/current/doku.php?id=fujitsu:hpcgateway:support:check_batch_manager)

Last update: **2018/01/17 19:12**

