

Troubleshooting LSF jobs

Why can't my job run now?

Once you submit your job to LSF using bsub, it enters the PENDING state. You can see all your pending jobs with

```
bjobs -p
```

You can see the status of a particular job with its JobID. Look for PENDING REASONS: in the output.

```
bjobs -p3 -l <jobid>
```

```
....
```

PENDING REASONS:

Job dependency condition not satisfied;

```
.....
```

It can be difficult to interpret the PENDING REASONS in the bjobs output. The cluster may just be very busy. You can see the cluster activity at <https://hpc-grafana.mskcc.org/>

Other LSF commands such as bhosts, lshost and lshost -gpu will give you current information about the available nodes and resources on the command line. You can also use RTM to view LSF details as the guest user at <http://lila-rtm01.mskcc.org/cacti/index.php> and <http://juno-rtm01.mskcc.org/cacti/index.php>

Things to check for:

- Typos in your bsub command.
- Any requested GPU models exist. lshost -gpu will list them with the correct syntax.
- The requested memory requirement (-R rusage[mem=4]) is in GB (gigabytes) and is PER SLOT (-n) and not per job.
- Make sure that you are in the SLA (Service Level Agreement) for any nodes that you specifically request.
- Your job must be able to finish before any scheduled downtime reservation.

The more resources that you request, the longer it will take for LSF accumulate the resources to satisfy them. Jobs which request resources that the cluster does not possess will remain in the pending state indefinitely. The maximum walltime on lilac is 7 days and on Juno is 31 days. Jobs that are less than 6 hours can run on any node. But those longer than 6 hours can only run on nodes with SLA or on a subset of the shared nodes.

Examples of different pending reasons and how to check for them.

1. Requested CPUs are not available. The cluster is busy. The job has been submitted to the ls03 host but ls03 doesn't have 30 free slots.

```
bjobs -p3 -l ...
```

```
Job <1135202>, User <sveta>, Project <default>, Application <default>, Status <PEND>, Queue <cpuqueue>, Job Priority <12>, Command <sleep 200000>, Share group charged </sveta>, Esub <memlimit>
```

```
Fri May 14 15:32:55: Submitted from host <lilac-ln02>, CWD <$HOME>, 30 Task(s),
```

```
Specified Hosts <ls03>;
```

```
Fri May 14 15:32:55: Reserved <30> job slots on host(s) <30*ls03> for <30> tasks;
```

```
Fri May 21 05:11:04: Job will start no sooner than indicated time stamp;
```

PENDING REASONS:

Candidate host pending reasons (1 of 120 hosts):

Affinity resource requirement cannot be met because there are not enough processor units to satisfy the job affinity request: ls03;

Non-candidate host pending reasons (119 of 120 hosts):

Not specified in job submission: lx10, lx11, lx12, lx13, lx14, boson, lt01, lt02, lt03, lt04, lt05, lt06, lt07, lt08, lt09, ld01, ld02,

ld03, ld04, ld05, ld07, lg01, lg02, lg03, lg05, lg06, lt0, lt11, lt12, lt13, lt14, lt15, lt17, lt18, lt19, lp01,

lp03, lp05, lp06, lp07, ls01, lp35, ls05, ls06, ls07, lv01, ls08, ls09, lt20, ly01, lt21, ly02, lt22, ly04, ly05, ly06, ly07, ly08, ly09, lp10, lp11, lp12, lp14, li01, lp16,

lp17, lp18, ls11, ls12, ls13, ls14, ls15, ls16, ls17, ls18, lp20, lu01, lu02, lu03, lp23, lu04, lp24, lu05, lp25, lu06, lu07, lp27, lx01, lu08, lu09, lx03, lx04, lx05, lx07, lx08, lx09, lu10, lp30, lp31,

lp33, lp34;

Load information unavailable: lw01, lw02, ld06, lt16, lp04, lp09, ly03, ls10, lp19, lp26, lx06, lp32;

Closed by LSF administrator: lg04, ls02, ls04, lx02;

2. Requested RAM(memory) is not available. The job asked for 400GB of memory on 5 nodes, which is 2,000GB RAM total: bsub -n 5 -R "span [ptile=5]" -R "rusage[mem=400]"

>bjobs -p3 -l ...

Mon May 17 14:28:26: Submitted from host <lilac-ln02>, CWD <\${HOME}>, 5 Task(s),

Requested Resources < rusage[mem=400] span[ptile=5]>;

PENDING REASONS:

Candidate host pending reasons (92 of 120 hosts):

Resource limit defined on host(s) and/or host group has been reached (Resource: slots, Limit Name: limit11, Limit Value: 68): lx09, lx08, lx04, lx03, lx10, lx12, lx13, lt10, lt15, lt17, ls01,

lu06, ls05, ls07, lu05, ls08, lt20, lt02, lt05, lt07, lu04, ls11, ls12, lu03, lu02, ls14, ls15, ls18;

Job slot limit reached: lt09, lx07, lx05, lx11, lt11, lx14, lu08, lx01, lt01, lu07, ls03, ls06, ls09, lt21, lt03, lt22, lt04, lt08, ls13;

Job's requirements for resource reservation not satisfied (Resource: mem): boson, lp30, lg01, lg03, lg05, lg06, lt14, lp01, lp07, ly01, ly02, lp24, ly04, ly05, ly06, ly08, lp16, lp17;

Resource limit defined on host(s) and/or host group has been reached (Resource: mem, Limit Name: limit11, Limit Value: 95): lu10, lg02, lu09, lt18, ly07, ly09, ls10, lu01;

Resource limit defined on host(s) and/or host group has been reached (Resource: mem, Limit Name: limit12, Limit Value: 78): lp34, lp27, lp06, lp33, lp31, lp12, lp18, lp20;

Resource limit defined on host(s) and/or host group has been reached (Resource: mem, Limit Name: limit1212, Limit Value: 78): lp03, lp25, lp10, lp11, lp14, lp35, lp23;

Affinity resource requirement cannot be met because there are not enough processor units to satisfy the job affinity request: lt12, lt13, lt19, ls17;

Non-candidate host pending reasons (28 of 120 hosts):

....

RUNLIMIT

60.0 min

MEMLIMIT

400 G

3. Requested RAM(memory) doesn't exist on cluster per host. Requested memory is per slot not per job.

>bjobs -p3 -l ...

Tue Sep 8 16:05:23: Submitted from host ,

CWD <\${HOME}>, 4 Task(s), Requested Resources <rusage[mem=200]>;

PENDING REASONS: Candidate host pending reasons (99 of 123 hosts):

Resource limit defined on host(s) and/or host group has been reached (Resource: mem, Limit Name: limit11, Limit Value: 95): lt15, lt17, lt18, lt19, lx14....

Job's requirements for resource reservation not satisfied (Resource: mem): lx10, lx12, lx13, boson, lt05, lt08, lt09, lx08, lx07....

Host is reserved to honor SLA guarantees: lp34, lp32, lp01, lp03, lp05, lp06....

Non-candidate host pending reasons (24 of 123 hosts):

Not specified in job submission: li01, lv01...

Load information unavailable: lp21, ls18, lp26, ls10, lp09, lp08, lg05, ld06

Closed by LSF administrator: lu04, lu05, ls05, lx09, lw02, lw01;

MEMLIMIT 200 G

RESOURCE REQUIREMENT DETAILS: Combined: select[(healthy=1) && (type == local)] order[!-slots:-maxslots] rusa ge[mem=200.00] span [hosts=1] same[model] affinity[thread(1)*1]</rusage[mem=200]>

This job won't run on Lilac cluster, because it requested 4x200=800GB of RAM on the same host span[hosts=1]

There is no host with 800GB on Lilac in cpuqueue.

Please, check resources:

```
>lshosts HOST_NAME type model cpuf ncpus maxmem maxswp
```

```
ls01 X86_64 GTX1080 60.0 72 512G ..
```

To check limits on resources:

```
>bresource
```

```
Begin Limit NAME = limit11 QUEUES = cpuqueue PER_HOST = ls-gpu/ lt-gpu/ lg-gpu/ lu-gpu/ lx-gpu/ lw-gpu/ ly-gpu/ SLOTS = 68 MEM = 95% ngpus_physical = 0 End Limit
```

4. Requested GPUs are not available. Again the cluster is busy.

```
bjobs -p3 -l....
```

```
#BSUB -n 1;#BSUB -gpu 'num=1';#BSUB -R 'span[ptile=1] rusage[mem=30]';#BSUB -q gpuqueue;
```

PENDING REASONS:

Candidate host pending reasons (92 of 123 hosts):

Job's requirements for resource reservation not satisfied (Resource: ngpus_physical): lx12, lx14, lt01, lt02, lt03, lt04, lt05, lt07, lt08, lt09, lu10, lx05, lx04, lx03, lg02, lu08, lt12, lt19...

Affinity resource requirement cannot be met because there are not enough processor units to satisfy the job affinity request: lt10, lt11, lx10, lt13, lt14...

Host is reserved to honor SLA guarantees: lp31, lp01, lp03, lp04, lp05, lp06, lp07, lp27, lp33, lp30, lp34, lp25, lp24, lp35, lp10...

Non-candidate host pending reasons (32 of 123 hosts):

Job's resource requirements not satisfied: lu02, ls13, lv01, lg06, ld07, ld05. Load information unavailable: ls18, lp21, ls10, lp26, lp09, lp08, lg05, ld06 Closed by LSF administrator: lu04, lu05, ls05, lx09, lw02, lw01;

Not enough GPUs on the hosts: ly03, lx01, lx02, lx06, lx11;

ESTIMATION: Tue Sep 8 17:00:21:

Started simulation-based estimation; Tue Sep 8 17:00:39: Simulated job start time on host(s) <1*lt03>

This job is waiting for mainly gpu resources to be available in Lilac cluster.

Estimated start time is :

Simulated job start time on host(s) <1*lt03>

```
bhosts -l lt03 HOST lt03 STATUS CPUF JL/U MAX NJOBS RUN SSUSP USUSP RSV DISPATCH_WINDOW ok 60.00 - 72 4 4 0 0 0 -
```

```
CURRENT LOAD USED FOR SCHEDULING: r15s r1m r15m ut pg io ls it tmp swp mem slots ngpus
```

```
Total 0.0 0.0 0.0 8% 0.0 19 1 76 48.7G 0G 269G 68 4.0
```

```
Reserved 0.0 0.0 0.0 0% 0.0 0 0 0 0G 0G 22G - 0.0
```

```
ngpus_physical healthy gpu_shared_avg_ut gpu_shared_avg_mut
```

```
Total 0.0 1.0 44.0 1.0
```

```
Reserved 4.0 0.0 0.0 0.0
```

5. GPU type doesn't exist on cluster or there is a typo or syntax problem with the GPU in the request.

```
> bjobs -p3 -l ..
```

Tue Sep 8 16:01:28: Submitted from host , CWD <\${HOME}>,
Requested Resources <select[gpu_model0=='geforcegtx1000']>, Requested GPU;
PENDING REASONS: Candidate host pending reasons (0 of 123 hosts).
Non-candidate host pending reasons (123 of 123 hosts):
Job's resource requirements not satisfied: lp35, lx10, lx11, lx12, lx13, lx14, boson, lt01, lt02, lt03, lt04, lt05, lt06, lt07, lt08
Not specified in job submission: ld01, ld02, ld03, ld04, ld05, ld07, lv01, l i01, lila-sched01, lila-sched02;
Load information unavailable: ld06, lg05, lp08, lp09, ls10, ls18, lp21, lp26
Closed by LSF administrator: lw01, lw02, ls05, lu04, lu05, lx09;
RUNLIMIT 10.0 min</select[gpu_model0=='geforcegtx1000']>

This job won't run because the gpu_model0 is not correct in bsub and this resource is not available on Lilac cluster :

Candidate host : 0

Correct name is GeForceGTX1080

>lshosts -gpu

HOST_NAME gpu_id gpu_model gpu_driver gpu_factor numa_id ls01 0

GeForceGTX1080 440.33.01 6.1 0 1

GeForceGTX1080 440.33.01 6.1 0 2

GeForceGTX1080 440.33.01 6.1 1 3

GeForceGTX1080 440.33.01 6.1 1

6. Nodes are in system level reservation used for rolling upgrade or scheduled cluster level downtime.

bjobs -p3 -l ...

Job <1117724>, User <sveta>, Project <default>, Application <default>, Queue <cpuqueue>, Job Priority <12>, Command <sleep 200000>,
Esub <memlimit>

Wed May 12 16:36:49: Submitted from host <lilac-ln02>, CWD </etc/security/limits.d>, 2 Task(s), Specified Hosts <lx10>;

PENDING REASONS:

Candidate host pending reasons (1 of 120 hosts):

Not enough slots or resources for whole duration of the job: lx10;

Non-candidate host pending reasons (119 of 120 hosts):

Not specified in job submission: lp35, lx11, lx12, lx13, lx14, boson, lt01, lt02, lt03, lt04, lt05, lt06, lt07, lt08, lt09, ld01, ld02

.....

7. Nodes are reserved under SLA.

When will my job start to run?

bjobs -l <jobid>

Check for "ESTIMATION" in the output

Details forthcoming

Why did my job exit abnormally?

bhist -l JID

bhist -n 0 -l JID

Details forthcoming