

[illegible][illegible]

bsub

[illegible]

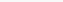
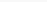
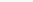
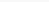
bsub

[illegible]

?????

```
bsub [options] command [arguments]
```

- [options] □ bsub □□□□□□□□ CPU□□
- command □□□□□□ MPI □□□□□□ mpirun □□
- [arguments] □□□□□□

“  e5v3ib  24 MPI 

```
$ bsub -q e5v3ib -n 24 "module load oneapi/2024.0/mpi && mpirun ./app"
Job <3206000> is submitted to queue <e5v3ib>
```

```
bsub < jobfile
```

jobfile  shell

```
#BSUB [options]
command [arguments]
```



bsub

```
f01n02 -m "f01n01+2
```



f01n02+1 f01n03" ☐

- -R "res\_req"

[illegible]

-R

largemem

- R "select[hname!=host\_name]" ☐ host\_name☐

&&☐☐☐☐

```
-R "select[hname!=x001 && hname!=x002]"  x001  x002
```

- X

- `-W [hour:]minute`  kill

CPU??

- `-R affinity[core:cpubind=core:membind=localprefer:distribute=pack]` 

[illegible][illegible]

?????

- $-r$

- `-Q "exit_code [exit_code ...]"` 

all [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ]      ~ [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ]

????

- $-I$

[illegible]

|  |  |  |  |
|--|--|--|--|
|  |  |  |  |
|--|--|--|--|

- K

- i input\_file

- `-o output_file` 

- `-e error_file` 

- %J   JOBID    -0  -00





```
-o /dev/null
```





```
-gpu [ ] [ ] [ ] [ ] [ ] [ ] : [ ] [ ] [ ] [ ] [ ] [ ] num=1:mode=shared:mps=no:j_exclusive=yes [ ] [ ] [ ] [ ] [ ] [ ]
```

- ????

[illegible]

- ```
• -w 'done(job_ID | "job_name")' [ ] job_ID [ ] job_name [ ] DONE  
[ ] 0  
• -w 'ended(job_ID | "job_name")' [ ] job_ID [ ] job_name [ ] EXIT [ ]  
DONE
```



- 

OpenMP (Open Multi-Processing)

MPI (Message Passing Interface)

OpenMP

MPI

Diagram illustrating four different memory access patterns for a 1D array of 20 elements, showing their compatibility with OpenMP, MPI, and NUMA:

- Pattern 1 (Sequential):** Accesses all 20 elements sequentially. Compatible with OpenMP, MPI, and NUMA.
- Pattern 2 (Strided):** Accesses elements 1, 5, 9, 13, and 17. Compatible with OpenMP and MPI.
- Pattern 3 (Random):** Accesses elements 1, 10, 19, 8, 17, 6, 15, 4, 13, 2, 11, 20, 9, 18, 7, 16, 3, and 12 in that order. Compatible with MPI and NUMA.
- Pattern 4 (Strided):** Accesses elements 1, 2, 3, 4, and 5 sequentially. Compatible with OpenMP and MPI.

Diagram illustrating the distribution of MPI ranks across different components:

- MPIRUN:** 10 MPI ranks (represented by small squares).
- LSB\_MCPU\_HOSTS:** 5 MPI ranks (represented by small squares).
- CPU:** 5 MPI ranks (represented by small squares).
- MPI/OpenMP:** 10 MPI ranks (represented by small squares).
- MPI:** 5 MPI ranks (represented by small squares).

The diagram shows that MPI ranks are distributed across mpirun, LSB\_MCPU\_HOSTS, CPU, and MPI/OpenMP. MPI/OpenMP ranks are distributed across MPI/OpenMP and MPI.

- ```
source /fs00/software/lsf/misc/ompthreads.sh [N]
```

3. ☐ MPI ☐ OpenMP ☐  
☐  
☐

OMP\_NUM\_THREADS

CPU

??????



?????

- LSB\_JOBID[ ] ID
- LSB\_QUEUE[ ]
- LSB\_JOBNAME[ ]
- LSB\_DJOB\_NUMPROC[ ] CPU[ ]
- LSB\_DJOB\_HOSTFILE[ ]
- LSB\_HOSTS[ ] CPU[ ]
- LSB\_MCPU\_HOSTS[ ] CPU[ ]

```
LSB_DJOB_NUMPROC=6
LSB_HOSTS="node1 node1 node1 node2 node2 node2"
LSB_MCPU_HOSTS="node1 3 node2 3"
```

```
$ cat $LSB_DJOB_HOSTFILE
node1
node1
node1
node2
node2
node2
```

```
“ LSB_HOSTS [ ] LSB_MCPU_HOSTS [ ] LSB_MCPU_HOSTS
[ ] LSB_HOSTS [ ] LSB_HOSTS [ ] 4096 [ ]
LSB_MCPU_HOSTS[ ]
```

??????

????

[ ] e52660 [ ]



```
$ bsub -q e52660 ./app
Job <3279929> is submitted to queue <e52660>.
$ cat job.lsf
#BSUB -q e52660
./app
```

```
$ bsub < job.lsf
Job <3279930> is submitted to queue <e52660>.
```

## MPI ????

**MPI**          **mpirun**  







```
bsub -q e5v4p100ib -gpu num=1 ./gpu_app
```

■■■■■

4■ GPU ■■■■

62v100ib ■■■■

GPU-CPU ■■

```
bsub -q 62v100ib -gpu "num=4:aff=yes" ./gpu_app
```

Revision #41

Created 9 May 2021 16:38:05 by Yao Ge

Updated 23 June 2024 16:06:54 by Yao Ge