

Envío e xestión dos traballos

Enviar os traballos ao sistema de colas

O envío de traballos realízase a través do comando `qsub`, cuxo argumento obrigatorio é o nome dun script de shell.

```
ct$ qsub script.sh
```

O comando `qsub` admite como parámetros as mesmas opcións que poden indicarse como comentarios `#PBS` no script.

Consultar o estado do traballo, as colas ou os nodos

Información das colas

O comando `qstat` permite consultar o estado das colas,

```
ct$ qstat -q # Información global das colas
server: ctcomp2
```

Queue	Memory	CPU	Time	Walltime	Node	Run	Que	Lm	State
graphic32	--	--	04:00:00	--	--	0	0	--	E R
np16	--	--	192:00:0	--	--	0	0	--	E R
np32	--	--	288:00:0	--	--	0	0	--	E R
especial	--	--	672:00:0	--	--	0	0	--	E R
parallel	--	--	192:00:0	--	--	0	0	--	E R
np2	--	--	192:00:0	--	--	0	0	--	E R
np8	--	--	192:00:0	--	--	0	0	--	E R
short	--	--	--	--	--	0	0	--	E R
graphic1	--	--	04:00:00	--	--	0	0	--	E R
np1	--	--	672:00:0	--	--	0	0	--	E R
batch	--	--	--	--	--	0	0	--	E R
np4	--	--	192:00:0	--	--	0	0	--	E R
interactive	--	--	01:00:00	--	--	0	0	--	E R
np64	--	--	384:00:0	--	--	0	0	--	E R
graphic	--	--	--	--	--	0	0	--	E R
bigmem	--	--	--	--	--	0	0	--	E R
graphic8	--	--	04:00:00	--	--	0	0	--	E R
						0	0		

A columna State indica coa súa primeira letra se a cola está (E)nabled ou (D)isabled e coa segunda letra se a cola está (R)unning ou (S)topped.

Información dos traballos

Cada vez que se envía un traballo asígnaselle un JOB_ID que serve como identificador único. Se o traballo enviouse coa opción -t entón identificarase mediante job_id[indice].

```
ct$ qstat # Información xeral dos traballos de usuario
Job id          Name          User          Time Use S
Queue
-----
--
999999.ctcomp2 nome_do_traballo nome_usuario 38:05:59 R np32
```

A columna Time Use amosa o tempo de CPU empregado. A columna S é o estado do traballo, que pode ser un dos seguintes:

- C - Job is completed after having run
- E - Job is exiting after having run.
- H - Job is held.
- Q - job is queued, eligible to run or routed.
- R - job is running.
- T - job is being moved to new location.
- W - job is waiting for its execution time(-a option) to be reached.

```
ct$ qstat -f 999999.ctcomp2 # Información sobre un traballo específico
Job Id: 999999.ctcomp2.innet
Job_Name = nome_do_traballo
Job_Owner = nome_usuario@ctcomp2.innet
job_state = Q
queue = np32
server = ctcomp2.innet
Checkpoint = u
ctime = Fri Feb 12 10:09:34 2016
Error_Path =
ctcomp2.innet:/home/local/nome_usuario/nome_do_traballo.e999999
Hold_Types = n
Join_Path = n
Keep_Files = n
Mail_Points = ae
Mail_Users = nome_usuario@usc.es
mtime = Fri Feb 12 10:09:34 2016
Output_Path =
ctcomp2.innet:/home/local/nome_usuario/nome_traballo.o999999
Priority = 0
qtime = Fri Feb 12 10:09:34 2016
Rerunable = True
Resource_List.nodect = 1:ppn=32:intel:xeonl
Resource_List.nodect = 1
```

```
Resource_List.nodes = 1:ppn=32:intel:xeonl
Resource_List.vmem = 63gb
Resource_List.walltime = 12:00:00
substate = 10
Variable_List = PBS_0_QUEUE=batch,PBS_0_HOME=/home/local/nome_usuario,
PBS_0_LOGNAME=nome_usuario,
PBS_0_PATH=/usr/local/bin:/usr/bin:/bin:/usr/local/games:/usr/games,
PBS_0_MAIL=/var/mail/nome_usuario,PBS_0_SHELL=/bin/bash,
PBS_0_LANG=es_ES.UTF-8,PBS_0_WORKDIR=/home/local/nome_usuario,
PBS_0_HOST=ctcomp2.innet,PBS_0_SERVER=ctcomp2
euser = nome_usuario
egroup = citius
queue_rank = 2110
queue_type = E
etime = Fri Feb 12 10:09:34 2016
submit_args = script.sh
fault_tolerant = False
job_radix = 0
submit_host = ctcomp2.innet
```

Unha característica interesante dos traballos rematados é o EXIT_STATUS que se amosaría cando el JOB_STATE es C.

Código interno	Valor de EXIT_STATUS	Significado
JOB_EXEC_OVERLIMIT	-10	
JOB_EXEC_STDOUTFAIL	-9	
JOB_EXEC_CMDFAIL	-8	Exec() of user command failed
JOB_EXEC_BADRESRT	-7	Job restart failed
JOB_EXEC_INITRMG	-6	Job aborted on MOM init, chkpt, ok migrate
JOB_EXEC_INITRST	-5	Job aborted on MOM init, chkpt, no migrate
JOB_EXEC_INITABT	-4	Job aborted on MOM initialization
JOB_EXEC_RETRY	-3	Job execution failed, do retry
JOB_EXEC_FAIL2	-2	Job execution failed, after files, no retry
JOB_EXEC_FAIL1	-1	Job execution failed, before files, no retry
JOB_EXEC_OK	0	Job execution successful
	1-256	Exit status of the top-level shell
	>256	Traballo rematado por unha sinal UNIX, restarlle 256 dáños o número do sinal.

```
ct$ checkjob 999999.ctcomp2 # Información sobre un traballo
específico

checking job 999999

State: Running
Creds: user:nombre_usuario group:citius class:np32 qos:DEFAULT
WallTime: 00:25:46 of 12:00:00
SubmitTime: Tue Feb 16 10:40:31
(Time Queued Total: 00:00:01 Eligible: 00:00:01)
```

StartTime: Tue Feb 16 10:40:32

Total Tasks: 32

Req[0] TaskCount: 32 Partition: DEFAULT

Network: [NONE] Memory >= 0 Disk >= 0 Swap >= 0

Opsys: [NONE] Arch: [NONE] Features: [active][intel][xeon]

Allocated Nodes:

[inode15:32]

IWD: [NONE] Executable: [NONE]

Bypass: 0 StartCount: 1

PartitionMask: [ALL]

Flags: RESTARTABLE

Reservation '137092' (-00:25:32 -> 11:34:28 Duration: 12:00:00)

PE: 32.00 StartPriority: 21

```
ct$ tracejob -n 3 999999.ctcomp2 # Devolve o contido dos logs relativos ao
jobid indicado.
```

Job: 136553.ctcomp2.innet

02/10/2016 15:22:26 S enqueueing into batch, state 1 hop 1

02/10/2016 15:22:26 S dequeuing from batch, state QUEUED

02/10/2016 15:22:26 S enqueueing into np1, state 1 hop 1

02/10/2016 15:22:26 S Job Run at request of citiuscap@ctcomp2.innet

02/10/2016 15:22:26 S Not sending email: User does not want mail of this type.

02/10/2016 15:22:26 A queue=batch

02/10/2016 15:22:26 A queue=np1

02/10/2016 15:22:26 A user=nome_usuario group=citius

jobname=nome_traballo queue=np1 ctime=1455114146

qtime=1455114146 etime=1455114146 start=1455114146

owner=nombre_usuario@ctcomp2.innet exec_host=inode19/24

Resource_List.nodect=1:ppn=1 Resource_List.nodect=1

Resource_List.nodes=1:ppn=1 Resource_List.vmem=2040mb

Resource_List.walltime=12:00:00

02/10/2016 16:08:34 S Exit_status=0 resources_used.cput=00:46:14

resources_used.mem=234868kb

resources_used.vmem=1002480kb

resources_used.walltime=00:46:08

02/10/2016 16:08:34 S on_job_exit valid pjob: 999999.ctcomp2.innet (substate=50)

02/10/2016 16:08:34 A user=nome_usuario group=citius

jobname=nome_traballo queue=np1 ctime=1455114146 qtime=1455114146

etime=1455114146 start=1455114146 owner=nombre_usuario@ctcomp2.innet

exec_host=inode19/24 Resource_List.nodect=1:ppn=1 Resource_List.nodect=1

Resource_List.nodes=1:ppn=1 Resource_List.vmem=2040mb

Resource_List.walltime=12:00:00 session=7304 end=1455116914 Exit_status=0

resources_used.cput=00:46:14 resources_used.mem=234868kb

```
resources_used.vmem=1002480kb resources_used.walltime=00:46:08
02/10/2016 17:08:35 S dequeuing from np1, state COMPLETE
```

Información dos nodos

Para obter unha vista global do estado do clúster, pódese empregar o comando nodes - usage.

```
$ nodes - usage
+-----+-----+
| USAGE                                | NODE                                |
+-----+-----+
| #####                                | node1 (64/64)                       |
| #####                                | node2 (64/64)                       |
|                                       | node3 (0/64)                         |
| #####                                | node4 (64/64)                       |
|                                       | node5 (0/64)                         |
| #####                                | node6 (64/64)                       |
|                                       | node7 (0/64)                         |
|                                       | inode11 (0/32)                      |
|                                       | inode12 (0/??)                      |
|                                       | inode13 (0/32)                      |
|                                       | inode14 (0/32)                      |
|                                       | inode15 (0/??)                      |
|                                       | inode16 (0/32)                      |
|                                       | inode17 (0/??)                      |
|                                       | inode18 (0/??)                      |
| ##                                   | inode19 (2/32)                      |
| #####                                | inode20 (28/32)                     |
+-----+-----+
| #####                                | TOTAL (286/640)                     |
+-----+-----+
```

Para obter información sobre os usuarios que se atopan en cada nodo, pódese empregar o comando node-users <nodo>:

```
$ node-users node1
Tracing node
jobs.....
jorge.suarez natalia.fernandez
```

Para obter información máis detallada sobre os nodos, pódese empregar o comando pnbsnodes:

```
ct$ pnbsnodes #Información detallada de todos os nodos
node1
state = free
np = 64
properties = amd,bigmem,test,active,activeX
ntype = cluster
status =
```

```
rectime=1455267717, varattr=, jobs=, state=free, netload=86957182662, gres=, loadave=0.00, ncpus=64, physmem=132250896kb, availmem=162914704kb, totmem=163499276kb, idletime=1876325, nusers=0, nsessions=0, uname=Linux node1 3.2.0-4-amd64 #1 SMP Debian 3.2.51-1 x86_64, opsys=linux
mom_service_port = 15002
mom_manager_port = 15003
gpus = 0

node2
state = down,offline
np = 64
properties = amd,bigmem
ntype = cluster
status =
rectime=1454919087, varattr=, jobs=, state=free, netload=1185896, gres=, loadave=0.00, ncpus=64, physmem=264633540kb, availmem=295220244kb, totmem=295881920kb, idletime=11140, nusers=0, nsessions=0, uname=Linux node2 3.2.0-4-amd64 #1 SMP Debian 3.2.51-1 x86_64, opsys=linux
mom_service_port = 15002
mom_manager_port = 15003
gpus = 0

.....
```

```
ct$ pbsnodes -l # Listaxe dos nodos apagados(down) ou non
dispoñibles(offline)
node2          down,offline
node3          down,offline
node4          down,offline
node5          down,offline
node6          down,offline
node7          down,offline
inode11        down,offline
inode12        down,offline
inode13        down,offline
inode14        down,offline
inode15        down,offline
inode17        down,offline
inode18        down,offline
inode19        down,offline
```

Eliminar un traballo da cola

O comando `qdel` permite ao usuario eliminar un traballo. Funciona enviándolle primeiro un sinal `TERM` e despois unha `KILL`. Este comando precisa como argumento o identificador que PBS asígnalle cando se rexistra un novo traballo, e que pode consultarse con `qstat`.

```
ct$ qdel job_id
```

From:

<https://wiki.citius.usc.es/> - Wiki do CiTIUS

Permanent link:

https://wiki.citius.usc.es/centro:servizos:hpc:envio_trabajo

Last update: **2017/10/09 11:29**

