

What if I want to lock down certain hosts to only run certain jobs?

There are several ways to do this in Qube, but a best practice that we recommend is to use the worker's cluster and restrictions as outlined in [How to use clustering for workers](#).

Essentially, you set a worker's [worker_cluster](#) and [worker_restrictions](#) to the same value, so that a worker:

- is in a particular cluster with [worker_cluster](#)
- will only accept job that are submitted into that cluster with [worker_restrictions](#)

Then, your users will submit jobs to that cluster by setting that job's **cluster** value to the name of the cluster; this job will have best priority on the workers in that cluster, but will be free to also run on hosts in other clusters (at a reduced priority).

If the job should **only** run on workers in that cluster, the user should also set the job's **restrictions** to the cluster name in which they wish the job to run.

Here's how I'd set up the [qbwrk.conf](#) to have hosts A - D all be in the `/project1` cluster, but hostC will **only** run jobs from that cluster.

```
# define the template, currently only specifies the cluster, but can be any
combination of parameters
[project1]
worker_cluster = "/project1"

# all 4 hosts inherit from the "project1" template
[hostA] : project1

[hostB] : project1

# hostC inherits the "project1" template, but also sets the
worker_restriction
[hostC] : project1
worker_restrictions = "/project1

[hostD] : project1
```

See Also

[Controlling Host Selection](#)