

__smartShare

This version of this page should not be displayed - please help us maintain our documentation and send an email to support@pipelinefx.com with the URL in which this page is appearing.

Overview

SmartShare is a feature in which the system automatically expands job instances to fill up idle, or "surplus", worker job slots. The expansion is automatically balanced among jobs, so that each job gets an equal number of surplus job slots.

Jobs have a 'max_cpus' value which puts an upper limit on how wide a single job can expand. This value can be supplied via a site-wide default or specified on a job-by-job basis.

SmartShare from a user's perspective

When SmartShare is enabled, jobs will automatically start using it. Note that SmartShare is only applicable to agenda-based jobs.

It is possible to control how much a job automatically expands, on a job-by-job basis, by setting a new job attribute called "max_cpus" at submission. This value sets an upper limit to the number of instances to which the job will expand. The max_cpus parameter is only visible in the submission UI when in Expert Mode. If an end-user does nothing, their job will expand as widely as the system administrator has defined (see Configuring SmartShare, below). The administrator can specify a default and a maximum value for max_cpus.

Jobs will expand to their max_cpus value if there are enough empty slots on the farm. As other users begin to submit jobs as well, a user's job will shrink to allow for a more "balanced" usage of the farm between jobs. No job will shrink below its "instance count" (or "CPUs") value. For example, a job submitted from the command line "qbsub" utility as

```
$ qbsub -priority 5 -cpus 10 -max_cpus 50 -range 1-100 sleep 60
```

ends up with 10 primary instances as specified by the "cpus" option at priority 5, and may expand up to 50 instances if surplus job slots are available. In other words, as many as 40 secondary instances will be dispatched, at a lower priority.

Settings

- Setting max_cpus to "*" (asterisk) or a negative number, such as "-1", means "no limit", which will allow the job to expand to the [supervisor_max_cpus_limit](#) setting. This is configured by the site administrator.
- Setting max_cpus to 0 disables SmartShare for the job.
- If max_cpus is unspecified at submission, the system will apply the default value, which is [supervisor_default_max_cpus](#). This is configured by the site administrator.
- The max_cpus value may also be modified after submission, in the GUI, with the [qbmodify](#) command, or via the API.

Priorities

- The instances that were automatically expanded by SmartShare ("secondary instances") run at low priority, and are the first subject to pre-emption in order to run primary instances from other jobs. Secondary instances may also be pre-empted in order to run secondary instances from other jobs, to balance the number of secondary instances among all jobs.
- The method of pre-emption can either be "aggressive", meaning the instance is immediately terminated, or "passive", meaning the instance is allowed to finish processing the current agenda item before terminating, and is configured by the site administrator (see [supervisor_smart_share_preempt_policy](#) below).
- The default smart share pre-emption mode is "aggressive"; secondary instances will be stopped immediately in order to accomplish the balancing of jobs as quickly as possible

Configuring SmartShare



SmartShare is automatically enabled for new 6.6+ installs and upgrades alike.

To opt out, the site administrator can set the [supervisor_smart_share_mode](#) to "none" in qb.conf on the Supervisor machine:

```
supervisor_smart_share_mode = none
```

-
- **supervisor_smart_share_mode**
 - Can be either "jobs" or "none".
 - default: "jobs"
 - Setting it to "jobs" will make the system try to balance the number of the automatically-expanded "secondary" instances on a job-to-job basis
 - Setting it to "none" disables SmartShare.
 - **supervisor_smart_share_preempt_policy**
 - method of preemption for secondary instances
 - aggressive, passive, disabled
 - default: aggressive
 - **supervisor_default_max_cpus**
 - used to set the site-wide default value of "max_cpus" if not given at submission
 - default: -1
 - Note: "*" can be used in the qb.conf file too, instead of -1
 - **supervisor_max_cpus_limit**
 - largest value of "max_cpus" that can be used on any job
 - default 100
 - Note: if any job specifies a value greater than this, the system will silently cap it to this value