Agents and Workers

Table of contents

Refer to The Grinder processes (../g3/getting-started.html#The+Grinder+processes) for an overview of the various processes. This page provides some further details.

1 Agent processes

It is typical to run a single agent process on each load injector machine.

When an agent is started, it attempts to connect to the console (../g2/console.html) . If it can connect, it will wait for a signal from the console before starting worker processes. Otherwise, the agent process will start a number of worker processes as specified by its local grinder.properties (../g3/properties.html) file.

If the network connection between the agent and the console is terminated, or the console exits, the agent will exit. If you want the agent to keep running and try regularly to reconnect to the console, use the `-daemon` command line switch. This might prove useful if you register an agent as an operating system service.

1.1 Summary of agent process options

Most agent options are controlled by the grinder.properties (../g3/properties.html) file. You can set properties on the command line (../g3/properties.html#Specifying+properties +on+the+command+line) .

`-daemon [reconnect time]`	If this option is specified on the agent command line, and the connection to the console cannot be established or the connection is lost, the agent will sleep for a while and then attempt to connect to the console again. The default sleep time is 60 seconds, but this can be controlled by providing a *reconnect time* in seconds.

2 Worker processes

Worker processes are started by a controlling agent process. The agent process passes each worker a set of properties (../g3/properties.html) that control its behaviour.