

Job Control Guide

Schrödinger Software Release
2015-2

Job Control Guide Copyright © 2015 Schrödinger, LLC. All rights reserved.

While care has been taken in the preparation of this publication, Schrödinger assumes no responsibility for errors or omissions, or for damages resulting from the use of the information contained herein.

Canvas, CombiGlide, ConfGen, Epik, Glide, Impact, Jaguar, Liaison, LigPrep, Maestro, Phase, Prime, PrimeX, QikProp, QikFit, QikSim, QSite, SiteMap, Strike, and WaterMap are trademarks of Schrödinger, LLC. Schrödinger, BioLuminate, and MacroModel are registered trademarks of Schrödinger, LLC. MCPRO is a trademark of William L. Jorgensen. DESMOND is a trademark of D. E. Shaw Research, LLC. Desmond is used with the permission of D. E. Shaw Research. All rights reserved. This publication may contain the trademarks of other companies.

Schrödinger software includes software and libraries provided by third parties. For details of the copyrights, and terms and conditions associated with such included third party software, use your browser to open [third_party_legal.html](#), which is in the docs folder of your Schrödinger software installation.

This publication may refer to other third party software not included in or with Schrödinger software ("such other third party software"), and provide links to third party Web sites ("linked sites"). References to such other third party software or linked sites do not constitute an endorsement by Schrödinger, LLC or its affiliates. Use of such other third party software and linked sites may be subject to third party license agreements and fees. Schrödinger, LLC and its affiliates have no responsibility or liability, directly or indirectly, for such other third party software and linked sites, or for damage resulting from the use thereof. Any warranties that we make regarding Schrödinger products and services do not apply to such other third party software or linked sites, or to the interaction between, or interoperability of, Schrödinger products and services and such other third party software.

May 2015

Contents

Document Conventions	v
Chapter 1: Introduction	1
1.1 Notes	2
Chapter 2: Running Jobs	3
2.1 The Job Life Cycle	3
2.2 Running Jobs From Maestro	5
2.2.1 Output Settings	6
2.2.2 Job Settings	6
2.2.2.1 Job Name	6
2.2.2.2 Host, Job Distribution, and Processors	8
2.2.2.3 Other Settings	9
2.3 Running Jobs From the Command Line	9
2.3.1 Running Scripts	12
2.3.2 The HOST, DRIVERHOST, and SUBHOST Options	13
2.3.3 The WAIT option	14
2.3.4 The LOCAL Option	14
2.3.5 Location of the Hosts File	15
2.4 Location of the Scratch Directory	15
2.5 Software Version Selection	17
2.6 License File Location	17
2.7 Environment Variables	17
2.8 Input and Output Files	18
2.9 Incorporation of Job Output	19
Chapter 3: Managing Jobs	21
3.1 The Job Database	21
3.1.1 The Job Record	22
3.1.2 Job Status	23

3.2 Managing Jobs From Maestro	25
3.2.1 The Jobs Table	26
3.2.2 The Details Tab	27
3.2.3 The File Tab	28
3.3 Managing Jobs From the Command Line	28
3.3.1 General Job Control Queries	29
3.3.2 Recovering Stranded Jobs	30
3.3.3 Using jserver	30
3.3.4 Purging the Job Database	31
3.3.5 Performing Actions on Jobs with jnanny	31
Appendix A: The postmortem Utility	33
Appendix B: Environment Variables	37
Appendix C: Secure File Transfer to Queue Hosts	43
C.1 Controlling Secure File Transfer	43
C.2 Altering Secure Zone Settings	44
C.3 Verifying Secure Transfer	44
Getting Help	47
Glossary	51
Index	53

Document Conventions

In addition to the use of italics for names of documents, the font conventions that are used in this document are summarized in the table below.

Font	Example	Use
Sans serif	Project Table	Names of GUI features, such as panels, menus, menu items, buttons, and labels
Monospace	<code>\$SCHRODINGER/maestro</code>	File names, directory names, commands, environment variables, command input and output
Italic	<i>filename</i>	Text that the user must replace with a value
Sans serif uppercase	CTRL+H	Keyboard keys

Links to other locations in the current document or to other PDF documents are colored like this: [Document Conventions](#).

In descriptions of command syntax, the following UNIX conventions are used: braces { } enclose a choice of required items, square brackets [] enclose optional items, and the bar symbol | separates items in a list from which one item must be chosen. Lines of command syntax that wrap should be interpreted as a single command.

File name, path, and environment variable syntax is generally given with the UNIX conventions. To obtain the Windows conventions, replace the forward slash / with the backslash \ in path or directory names, and replace the \$ at the beginning of an environment variable with a % at each end. For example, `$SCHRODINGER/maestro` becomes `%SCHRODINGER%\maestro`.

Keyboard references are given in the Windows convention by default, with Mac equivalents in parentheses, for example CTRL+H (⌘H). Where Mac equivalents are not given, COMMAND should be read in place of CTRL. The convention CTRL-H is not used.

In this document, to *type* text means to type the required text in the specified location, and to *enter* text means to type the required text, then press the ENTER key.

References to literature sources are given in square brackets, like this: [10].

Introduction

Nearly all computational jobs launched from Maestro are run under Schrödinger's Job Control facility. The Job Control facility provides a uniform mechanism for launching, monitoring and controlling calculations, both for jobs launched from Maestro and for jobs launched from the command line. The Job Control facility keeps information on jobs in a database that is set up for each user.

The provision of a Job Control facility makes it easy to run Schrödinger software. You can submit jobs *from* any computer *to* any computer or cluster, or to a batch queue, and have the results returned to the computer from which you submitted the job. The process is independent of the particular platform from which you submit the job or on which you run the job. The platform details are hidden, but that also means that the details have to be communicated to Job Control.

The setup and configuration tasks necessary for Job Control to work are described in [Chapter 7](#) of the *Installation Guide*. In this manual, [Chapter 2](#) describes the life cycle of a job and provides information on submitting jobs and how the job execution is set up. [Chapter 3](#) provides information on the job database and managing jobs.

Schrödinger software was originally designed to run on hosts that run some version of the UNIX operating system (including Linux and vendor-specific implementations of UNIX). On Linux, you can run jobs and Job Control commands from the command line. On Windows, many of the programs and commands can be run directly from a Windows shell. Two shells are provided with the installation that have the Schrödinger environment set up:

- Schrödinger Command Prompt—DOS shell.
- Schrödinger Power Shell—Windows Power Shell (if available).

You can open these shells from Start > All Programs > Schrodinger-2015. If you want access to Unix-style utilities (such as `awk`, `grep`, and `sed`) or commands that are not available in these shells, simply preface the commands with `sh`, or type `sh` in either of these shells to start a Unix-style shell.

1.1 Notes

- Job Control changes made for Suite 2011 are not compatible with earlier releases. You must ensure that the hosts file does not point to software installations for earlier releases, and that the location of the job database is different from that for earlier releases (this is the case by default, if you do not explicitly set the location of the job database).

Running Jobs

Jobs can be submitted to a designated host from a Maestro session or from the command line. The details of setting up the data in Maestro or submitting jobs from the command line for a given product are described in the documentation for the product. However, there are some common elements that relate to the job submission and execution, and the incorporation of results into a Maestro project. These common elements are described in two sections of this chapter.

When you run a Schrödinger job on a particular host, Job Control determines where to find the hosts file, which version of the software to use, which environment variables need to be passed to the host, where the scratch directory is located, how to make the input files available to the job, how to retrieve the output files, and how to incorporate them into a Maestro project. An overview of this process is provided in the first section of this chapter. The remaining sections provide details on how Job Control obtains information and sets up the runtime environment for the job.

2.1 The Job Life Cycle

To understand how information is obtained and passed on by Job Control, it is useful to have an understanding of how Schrödinger jobs are run. The “life cycle” of a job can be summarized as follows.

1. A top-level script is run locally (on the *submission host*). This script parses command-line arguments relating to Job Control and sets some environment variables.
2. The top-level script then runs the startup script for the program locally. This script parses the command arguments for the program, and assembles information on input and output files for the program.
3. The startup script then runs a job launch script locally, called `jlaunch`. This script assigns a job ID, creates a record in the job database, and populates it with initial values. It also verifies that the required input files exist.
4. If the job is to be run on a remote host, `jlaunch` starts a script locally, called `jserver`, that sets up a socket connection to the remote host that can be used for file transfer. For local jobs, `jserver` is also started, but it does not need to manage file transfer, only the interaction with the job database.

5. The job launch script then runs a script on the execution host to start the actual program. This script is called `jmonitor`. If the job is remote, it copies input files to the execution host using the connection established by `jserver`, runs the program, and copies output files back to the submission host. The program itself is responsible for checking in and checking out licenses.
6. The job launch script takes care of cleanup and incorporation into a Maestro project, if required, and updates the job record with the final status.

If you submit a job to a batch queue, there are some extra tasks that need to be performed. The life cycle of a batch job is as follows.

1. A top-level script is run locally (on the *submission host*). This script parses command-line arguments relating to Job Control and sets some environment variables.
2. The top-level script then runs the startup script for the program locally. This script parses the command arguments for the program, and assembles information on input and output files for the program.
3. The startup script then runs a job launch script locally, called `jlaunch`. This script creates a record in the job database and populates it with initial values. It also creates a batch script for the job and submits this script to the queue.
4. The job launch script starts a script locally, called `jserver`, that sets up a connection to the queue manager that can be used for file transfer. By default, this connection is an SSH tunnel, but if secure transfer is not in effect, a socket connection is used. Another `jserver` process is started on the queue manager (or other designated host) that sets up connections to the compute nodes when the queued jobs are started. This process acts as a proxy for the file transfer, and is done because it is not always possible to connect directly from the submission host to the compute nodes. This instance of `jserver` is also known as a `jproxy` process.
5. When the queued job is started on an execution host (a cluster node, for example), it runs `jmonitor`. `jmonitor` then copies input files to the execution host, runs the program (which checks out and checks in licenses), and copies output files back to the submission host. The copying is done via the SSH tunnel or the sockets opened by `jserver` processes on the manager node and the submission host. If the queue has license checking, the queueing software checks for license availability. See [Section 7.9](#) of the *Installation Guide* for information on setting up license checking.
6. The job launch script takes care of cleanup and incorporation into a Maestro project, if required, and updates the job record with the final status.

Jobs that are distributed over multiple processors usually have a script that manages the distribution of subjobs to the individual processors. Each subjob is then executed by `jmonitor`.

2.2 Running Jobs From Maestro

To run a job from most Maestro panels, you use the Job toolbar.



You can start a job immediately by clicking Run. The job is run on the currently selected host with the current job settings and the job name in the Job name text box. If you want to change the job name, you can edit it in the text box before starting the job. Details of the main job settings are reported in the status bar, which is below the Job toolbar.

To make settings for the job, you can open the Job Settings dialog box by choosing Job Settings from the Settings button menu.



The dialog box is usually named *task* – Job Settings, where *task* is a name related to the product, job, or initiating panel. The controls available in the Job Settings dialog box vary according to the job being run. (Some products and solutions have separate locations in which job settings are made.)

When you start Maestro, the settings in the `schrodinger.hosts` file are used to determine the available host-related options in the Job Settings dialog box. Maestro searches the following directories for a `schrodinger.hosts` file, in the order given, and uses the first one that it finds.

- The Maestro startup directory, which is the directory in which you started Maestro (Linux) or the Schrodinger folder in your documents folder (Windows or Mac)
- The Schrödinger user resources directory, `$HOME/.schrodinger` (Linux or Mac) or `%USERPROFILE%\Schrodinger` (Windows)
- The installation directory

From the command line you can always determine which `schrodinger.hosts` file is being used by entering the following command in the Maestro startup directory:

```
[${SCHRODINGER}/]program -HOSTS
```

where *program* is any Schrödinger executable. The options in the Job Settings dialog box are passed on to Job Control.

The Job Settings dialog box can have two sections, Output and Job. The components available in each section are described in the sections below. Some of the components are present in all instances of this dialog box, others are only present in some instances. Examples of Job Settings dialog boxes are shown in [Figure 2.2](#).

Once you have finished making settings, you can click **Run** to start the job, or click **OK** to save settings without starting the job. This allows you to make job settings before making calculation settings, for example. Settings are saved for each application, with a few panels that have separate settings (MacroModel Start, PrimeX Calculate R Factors).

2.2.1 Output Settings

The item described below is present in most instances of the dialog box. Some instances have custom features, which are described in the application documentation.

- **Incorporate option menu**—Choose whether the new entries are appended to the project either individually or as an entry group, replace the existing entries, or are not incorporated into the project at all. See [Section 2.9 on page 19](#) for more information on incorporating output.

The choice that you make from this menu is persistent for a given job type: the next time you run a job of that type, the incorporation mode that you last used is the default mode. The incorporation mode is stored as a preference by Maestro, so the choice persists across Maestro sessions.

2.2.2 Job Settings

The **Job** section of the **Job Settings** dialog box contains tools for setting the job name, selecting the host, distributing the work over multiple processors, and a few other settings.

2.2.2.1 Job Name

At the top of this section are the job name controls, which vary according to the application.

- **Name text box**—Enter a name for the computational job. This name is used as the first part of the file name for the files created for the job, and also for the job directory if the application creates a directory for the job in the current working directory. A default name is assigned based on the job being run. Job names cannot contain spaces or non-printing characters.

The initial name shown is a standard name for the application, which might include calculation settings. A standard name that contains settings is updated when the settings change. You can modify a standard name; if you do and it still retains some settings, the name is updated if the settings change. You can also replace the entire job name.

When a standard name or modified standard name is used, the job name is made unique by appending an integer. This is done by checking for job directories or files in the current working directory. However, if you replace the job name to create a custom job name, the name might no longer be unique, and it is not automatically made unique. In that case, a warning is posted before any files are overwritten.

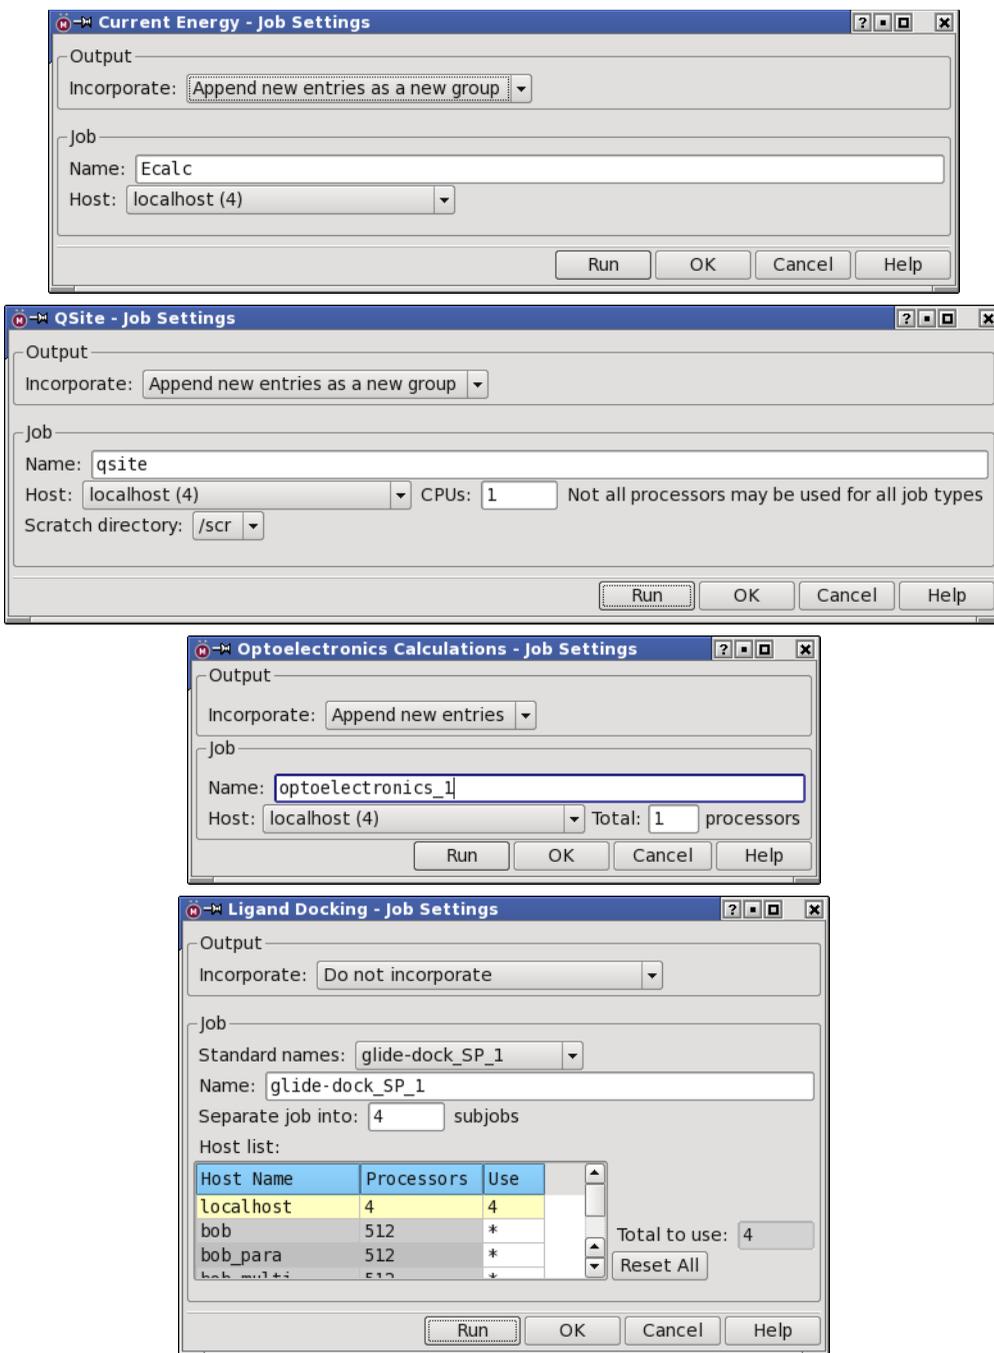


Figure 2.2. The Job Settings dialog box with various controls.

After a job is submitted, a new job name is automatically created for the next job from the current job name, by appending an integer or incrementing the integer. This is done for custom job names as well as standard or modified names.

2.2.2.2 Host, Job Distribution, and Processors

Below the job name settings are the host and job distribution settings. There are four main groupings. Jobs that cannot be distributed usually only have a hosts option menu. Jobs that can be divided into a specified number of independent subjobs and distributed over multiple hosts have a text box for the number of subjobs, and a host list for the hosts. Jobs that can be run in parallel or can only be distributed on a single host have a host option menu and a CPUs text box. Jobs that can be run on GPUs have a GPGPU list as well as the host option menu.

As terminology has changed with the introduction of multi-core CPUs, in the descriptions below, “processor” is taken to mean a single processing unit, such as a core.

- **Host**—Choose a host on which to run the job. This option menu displays all the hosts defined in the hosts file, with the number of processors on the host in parentheses.

The host `localhost` means the host on which you are running Maestro. If you run a job locally, Maestro automatically reduces the priority of the job so that it does not compete with Maestro for resources. The exceptions are Hydrophobic/philic map jobs, and structure cleanup jobs. To change this behavior, set the `SCHRODINGER_NICE` environment variable (see [Table B.1 on page 37](#)).

- **CPUs or Total processors**—Specify the number of processors to use for the job. The number of processors actually used may be limited by the number of licenses available for the type of job, and for some jobs, the number of processors used might vary during the course of the job.
- **Distribute subjobs across**—Choose an option for how to distribute each subjob for parallel execution. These options apply to Jaguar jobs, which can use both MPI processes and OMP threads for parallelization.
 - **CPUs**—Specify only the number of processors to use. The number of MPI processes and threads are determined automatically for some job types, while for others, the subjobs are distributed over the processors, one per processor.
 - **MPI processes**—Specify the number of MPI processes and the number of threads per process to use, and the number of subjobs to run simultaneously. The total number of processors requested is the product of these three numbers, as each subjob runs with the specified number of MPI processes and threads. When this option is selected, text boxes for the numbers are displayed, and the total number of CPUs that result from the values in the text boxes is reported to the right.

- **Separate job into N subjobs**—Specify the number of subjobs into which the job is split. Some drivers optimize the number of subjobs, so the number actually used might not be exactly the number given. For good load balancing, the number of subjobs should be at least 3-4 times the number of processors over which the job is distributed. The number of subjobs should not be smaller than the number of processors: if it is, it is reset to the total number of processors. See *Running Distributed Schrödinger Jobs* for more information.
- **Host list**—Choose the hosts on which to run the jobs by selecting rows in the table, and specify the number of processors to use for each host by editing the Use column. You can select multiple hosts with shift-click and control-click. Only the selected hosts are used for the job. You cannot mix queue hosts and non-queue hosts, and only one queue host can be selected. The Processors column shows the number of processors available on each host. The Total to use text box shows the total number of processors chosen. This text box is noneditable, so you must use the table to specify the number of processors.
- **Reset All**—Reset the Host list table to the default values (use all processors on the local host).
- **Automatically select GPUs**—If a host is selected from the Hosts option menu that has GPU units available for use, this option becomes available. The Total processors text box is relabeled Total GPUs, and the number of GPUs selected automatically can be specified in this text box.
- **GPU table**—Choose the GPU units on which to run the job, by selecting the table rows. Click to select, click again to deselect. The table is only populated when you choose a host from the Host option menu that is identified as having GPU units available for use, and is only available if you deselect Automatically select GPUs. The number of GPUs you select is displayed in the Total GPUs text box, which is noneditable when selecting GPUs.

2.2.2.3 Other Settings

There are other settings that appear in some Job Settings dialog boxes:

- **Scratch directory**—Choose a directory to store temporary files.
- **Entry title text box**—Enter a title for the input structure file. The default is the entry title.

2.3 Running Jobs From the Command Line

For most purposes, you can start jobs from Maestro. You can also run or submit jobs from the command line on Linux hosts or from a Schrödinger shell on Windows hosts (see below). The Job Control facility recognizes a number of command-line options that can be used to control the behavior of the job. These options are summarized in [Table 2.1](#).

Table 2.1. Command options.

Option	Description
-DISP <i>policy</i>	Set the project incorporation policy for this job: ignore, append, appendungrouped, workspace, or replace. Requires -PROJ, and overrides SCHRODINGER_JOB_DISPOSITION. The policies are described in Section 2.9 on page 19. Default: ignore.
-DRIVERHOST <i>host</i>	For distributed jobs, specify the host on which to run the driver (“master job”). Must specify a single host.
-HOST <i>host</i> -HOST <i>host:n</i> -HOST " <i>host1:n1 host2 ...</i> "	Run a job on the specified host or submit a job to the specified batch queue. <i>host</i> is the value of a name entry in the hosts file or the actual address of a host. To specify multiple hosts, supply a blank-separated list in quotes. To specify the (maximum) number of subjobs to run on a host, append a colon and the number of subjobs to the host name. Default: run on the local host.
-NICE	Run the job at reduced priority. On Linux, the program is run with <code>nice -19</code> . On Windows, the program is run at “idle” priority. Overrides SCHRODINGER_NICE.
-NOLAUNCH	Perform all the steps necessary for launching the job, but stop short of actually launching it.
-NONICE	Do not run the job at reduced priority. Overrides SCHRODINGER_NICE.
-NOSTARTUP	In the process of launching the job, stop short of executing the application startup script. Mainly useful for debugging.
-PROJ <i>projectname</i>	Assign the job to a Maestro project. Overrides SCHRODINGER_PROJECT
-QARGS <i>queue-args</i>	Pass arguments to the queue manager. These arguments are appended to those specified by the <code>qargs</code> settings in the hosts file and the SCHRODINGER_QUEUE_ARGS environment variable.
-SAVE	Copy the archived contents of the job directory back to the submission directory after the job finishes, as <i>jobid-jobdir.zip</i> .
-SUBHOST <i>host</i> -SUBHOST <i>host:n</i> -SUBHOST " <i>host1:n1 host2 ...</i> "	For distributed jobs, specify the hosts on which the subjobs will run. The syntax is the same as for -HOST. Default: run subjobs on the hosts specified by -HOST.
-TMPDIR <i>directory</i>	Specify the scratch directory for the job. The job directory is created as a subdirectory of the scratch directory. Overrides any user setting of SCHRODINGER_TMPDIR
-USER <i>username</i>	Specify the user name to be used for remote jobs. Must be used with -HOST. Default: use the same user name as on the submission host.

In addition to these options, the startup scripts for some programs support several other options. These options are summarized in [Table 2.2](#). You should check which options are supported by entering the command

```
$SCHRODINGER/program -HELP
```

Table 2.2. Options supported by some programs

Option	Description
-LOCAL	Write temporary files in the submission directory instead of the scratch directory. Input and output files are not copied. Default: write temporary files in the scratch directory.
-NOJOBID	Run the job outside of Job Control. If you use this option, you are responsible for managing all the environment variables, file handling, and checking done by Job Control.
-NOLOCAL	Write temporary files in the scratch directory instead of the submission directory. Some programs and utilities write files locally when run on the local host; this option can be used to force the use of the scratch directory.
-WAIT	Wait for the job to finish before executing another command. In a terminal window, this means that the command prompt is not displayed until the job finishes. In a script, it means that the next command is not executed until the job finishes. Default: return control to the shell immediately.

Command-line options always take precedence over the corresponding environment variable. Some of the options from [Table 2.1](#) and [Table 2.2](#) are described in more detail below. You can also obtain information for each program about the hosts you can use. These options are listed in [Table 2.3](#).

Table 2.3. Information options.

Option	Description
-DEBUG	Show the details of operation of the top-level script.
-DDEBUG	Show the verbose details of operation of the top-level script.
-ENTRY	Show the section of the <code>schrodinger.hosts</code> file that will be used for this job.
-HELP	Display command syntax for the application.
-HOSTS	List the hosts that are available for calculations.

On Windows, you can run Unix (Linux) commands by opening a Schrödinger Command Prompt window from the Start menu, then entering the `sh` command. The `unxutils` package provides many of the commands available in a Unix shell.

2.3.1 Running Scripts

You can run scripts from the command line with the following command:

```
$SCHRODINGER/run scriptname [-FROM product] [-FIND|-FINDALL] [options]
```

This command runs a top-level script like those for Schrödinger products, which recognizes the Job Control options listed in the tables above (except for Table 2.2) and sets up environment variables. The script name can be one of the scripts supplied with the distribution or downloaded from the Script Center, or it can be one of your own scripts. The script will not run under Job Control unless it has been set up to do so. Any options that are not recognized by the top-level script are passed on to your script.

The `run` command looks for scripts in the following locations in the order given:

1. The current directory.
2. The directory defined by the `MMSHARE_EXEC` environment variable, which is set by the `run` command to `$SCHRODINGER/mmshare-vversion/bin/platform`, where `version` is the `mmshare` version. (You should not set this environment variable yourself.)
3. The directory defined by the `SCHRODINGER_SCRIPTS` environment variable.
4. The `suite2015-2/scripts` directory in your Schrödinger user resource area, which is `$HOME/.schrodinger` on Linux and Mac and `%APPDATA%\Schrodinger` on Windows. This is the directory used for all scripts that are downloaded for personal use from the web site using the Update Scripts From Website panel in Maestro.
5. The directory `$SCHRODINGER/mmshare-vversion/python/common`.
6. The directory `$SCHRODINGER/mmshare-vversion/python/scripts`.
7. Your execution path, defined by the `PATH` environment variable.

The `-FIND` option displays the path used to find the script, telling you which of these directories would be used to run the script. The `-FINDALL` option displays the path to all of the possible locations where the script could be found, in the order given in the list above. These options are useful when you have several versions of a script and want to find out which one is being used, or which versions you have available.

If you use the `-FROM product` option, the executable directory for the product that you specify, `$SCHRODINGER/product-vversion/bin/platform`, replaces `MMSHARE_EXEC` in the hierarchy.

The `run` command is especially useful for scripts that make use of the Schrödinger libraries or Python modules, for which the environment is set up when you use this command.

2.3.2 The `HOST`, `DRIVERHOST`, and `SUBHOST` Options

Jobs can be submitted to a remote host using the `-HOST` option to specify the remote host name. This name must be the value of a name setting in the hosts file (a “host entry name”), or an actual host name. For example:

```
$SCHRODINGER/bmin -HOST host jobname
```

For programs that can run a single job in parallel or distribute several jobs over a number of processors, the `-HOST` option can be used to specify the list of hosts to be used. The host list is a list of host entry names, separated by spaces. The list must be enclosed in quotes if there is more than one host specified.

Each host entry name can also specify how many subjobs to run (at most), using the syntax:

```
host:subjobs
```

An example of a host list specification is:

```
-HOST "florence:2 glinda"
```

The first host in the list is the main host for the job, that is, the host on which the driver process for the parallel or distributed job runs. The entire set of hosts is used for the subjobs, including the first. Some products run the driver on the local host, regardless of the `-HOST` value.

For other products, the host on which the driver runs and the host on which the subjobs run can be specified independently, with the `-DRIVERHOST` option and the `-SUBHOST` option. This is useful when you want the driver to run locally, for example. You might want to do this when running the subjobs on a cluster via a queueing system, so that the driver is not occupying a node on the cluster. This is not necessary for Phase, which runs a subjob on the same host as the driver, and ignores the driver host setting.

If you use either of these two options, the `-HOST` option serves as a fallback for the other option. So if you use `-SUBHOST`, `-HOST` only specifies the location of the driver; if you use `-DRIVERHOST`, `-HOST` only specifies the location of the subjobs. If you use both, `-HOST` is ignored.

The relationship between these options can be understood in terms of how the driver host and the subjob hosts are set. `-HOST` sets the driver host and the subjob hosts. `-DRIVERHOST` unconditionally sets the driver host (overriding `-HOST`), and `-SUBHOST` unconditionally sets the subjob hosts (overriding `-HOST`). The default for anything that is not set is `localhost`. This means that if you only set `-DRIVERHOST`, the entire job runs on the driver host.

Not all jobs use all of these settings. See *Running Distributed Schrödinger Jobs* for detailed information.

If you specify a host name that is not a host entry name (e.g. an IP address or fully qualified domain name), the settings for `localhost` in the hosts file are used for the job. You must ensure that there is a software installation on the specified host at the same path as the installation on the local host that you used to start the job.

2.3.3 The WAIT option

All jobs are run in the background automatically, dissociated from the terminal session or application from which they were launched. As a result, the job continues to run even if you quit Maestro or the terminal session from which you launched the job. For command-line jobs, this means that the UNIX command prompt is displayed immediately, without waiting for the job to finish. This behavior is not always desirable, especially if you want to run the job in a script, in which some subsequent action can be taken only after the job finishes. The command-line option `-WAIT` can be used to prevent the shell from continuing to the next command until after the job finishes. For example:

```
‡SCHRODINGER/bmin -WAIT job_name
```

Even with this option, however, the calculation is placed in the background, so pressing `CTRL+C` or `CTRL+Z` does not affect it. This option applies to jobs submitted to a batch queue as well as jobs that are run directly.

When the job finishes, the return code for the job is set to the return code for the computational program. Thus, if the program failed, the return code for the job indicates why it failed. The return code is also available in the `ExitCode` field of the job record.

2.3.4 The LOCAL Option

For some applications, you can request temporary files to be written to the submission directory (the directory from which you submitted the job), instead of to the job directory, which is a subdirectory of the scratch directory. This is done with the `-LOCAL` option. For example,

```
‡SCHRODINGER/bmin -LOCAL jobname
```

Not every application supports this option. When it is supported, it may be used for both local and remote jobs. If you run a remote job with the `-LOCAL` option, it is important to make sure that the submission directory is accessible on the remote machine. This option also suppresses the copying of input and output files.

For some utilities, temporary files are written in the submission directory if the job is run locally. For these utilities, you can force the files to be written in the scratch directory with the `-NOLOCAL` option.

2.3.5 Location of the Hosts File

When you start a job from the command line, one of the first tasks of Job Control is to locate a hosts file. The hosts file used for a given job is the first one found in the following list:

- The file specified by the environment variable `SCHRODINGER_HOSTS`
- The `schrodinger.hosts` file in the current directory
- The `schrodinger.hosts` file in `$HOME/.schrodinger` (Linux/Mac) or `%USERPROFILE%\Schrodinger` (Windows)
- The `schrodinger.hosts` file in the Schrödinger software installation

For MacroModel jobs the following locations are searched before the list given above:

1. A file specified by the command line argument `-HOSTFILE`.
2. The file `jobname.hst` in the startup directory on the submission host, where *jobname* is the stem of the command file name for the current calculation (e.g., if the command file were called `cal_en1.com`, this file would be called `cal_en1.hst`).

The information in the hosts file is then used to make default settings for the job, such as the location of the scratch directory.

If the hosts file is specified by the environment variable `SCHRODINGER_HOSTS`, this file is also used to locate the hosts for any subjobs.

2.4 Location of the Scratch Directory

Most jobs now run in a scratch directory by default, rather than in the directory from which the job was started (the *submission* directory, also called the *launch* directory). When a job runs in a scratch directory, a subdirectory is created in it for the job, named `tmpdir/username/unique_name`. Here, *tmpdir* is the path to the scratch directory. The job name is usually used for *unique_name*, but it can also have a sequence number or the job ID appended to it to make the directory name unique. This subdirectory is called the *job directory*. Input files are copied to the job directory, temporary files, log files, and output files are created in the job directory, and the output and log files are copied back to the submission directory when the job finishes.

If you submit a job from Maestro, you may be able to choose the scratch directory in the Start dialog box. Maestro reads all the `tmpdir` settings from the hosts file, and presents these in the Scratch directory option menu.

For jobs submitted from the command line, there are a number of ways to specify the *tmpdir* directory. Job Control uses the first specification found from the following list:

- The directory given on the command line with the `-TMPDIR` option. For example,

```
$SCHRODINGER/bmin -TMPDIR /scr/mmod_tmp job_name
```
- The directory specified by the `SCHRODINGER_TMPDIR` environment variable, if this is set on the submission host.
- The directory specified by the first `tmpdir` setting for the host entry in the hosts file.

For jobs run on a remote host, the following locations are considered after the ones listed above if a scratch directory is not defined:

- The directory specified by the environment variable `SCHRODINGER_TMPDIR` on the remote host.
- The directory specified by the environment variable `TMPDIR` on the remote host.

If no other specification for *tmpdir* is found, the directory `$HOME/.schrodinger/tmp` is used. In this case, the *username* is not used to form the job directory name, since it would be redundant. Temporary files can occupy a large amount of disk space, so you are strongly advised to ensure that a suitable directory is defined on every host, preferably in the hosts file.

In all cases, *tmpdir* is created if it does not exist. If the file system on which *tmpdir* is to be created does not exist, the job fails.

When the job finishes, the job directory is automatically removed, if the following conditions are met:

- The output files were all successfully copied back to the submission directory.
- The directory did not exist before the job started.
- The `-SAVE` option was not used in the job submission.

Some programs allow you to force the storage of temporary files in the submission directory, by using the `-LOCAL` option to the command for running the program—see [Section 2.3.4 on page 14](#).

2.5 Software Version Selection

When a job is started, the Job Control facility checks that the software installation specified for the job exists and that the software (mmshare) version matches that on the submission host.

The software installation used on the execution host is determined by the following rules:

1. The installation specified by the `schrodinger` setting in the entry for the execution host in the hosts file on the submission host is used, if it is defined. If the installation does not exist on the execution host, the job fails.
2. The installation specified by the `SCHRODINGER` environment variable on the submission host is used, if there is no `schrodinger` setting in the hosts file for the execution host.

Thus, all the information on the software location must be available on the submission host. See [Section 7.1](#) of the *Installation Guide* for information on setting up the hosts file.

If you submit a job to a batch queue, an additional check is done when the job starts, using the information obtained on the submission host, which is passed on by Job Control. This is done because Job Control can only check the availability of software on the queue manager when the job is launched.

2.6 License File Location

When an executable starts or the queuing software checks the availability of licenses, the following locations are searched, in order, for the license file; the first one found is used.

1. `$SCHROD_LICENSE_FILE`
2. `$LM_LICENSE_FILE`
3. `$SCHRODINGER/license` or `$SCHRODINGER/license.txt`.
4. On a Mac, `/Library/Application Support/Schrodinger/license[.txt]`.

2.7 Environment Variables

Environment variables that can be used on the execution host can come from the following sources:

- Environment variables set in the shell startup script (`.bashrc`, `.cshrc`, and so on)
- Environment variables set on the submission host when the job is launched.
- Environment variables set in the hosts file for the execution host or `localhost`.

The environment variables are passed to the job in the following order of precedence:

1. Settings in the host entry.
2. Settings in the user environment on the submission host.
3. Settings in the user environment on the execution host. These are used as a fallback if those environment variables were not defined and exported with the job.

2.8 Input and Output Files

When you run a job from the command line you can specify input files in either of two ways: using absolute paths or relative paths. The way Job Control treats these differs. The configuration considerations for output files and for updates of the Job Control database are similar to those for input files specified using absolute paths. The basic rules are as follows:

- **Files specified using absolute paths.** These are paths that begin with a forward slash (/) such as `/home/unohu/prime/1mcp.mae`. Job Control first looks for the file on the execution host exactly as given. For instance, it would look for a file named `/home/unohu/prime/1mcp.mae` on the execution host. If it finds the file, it does nothing more, and the file on the execution host is used for input. If it does not find the file, it looks for a file of that name on the submission host, and copies it to the job directory of the execution host.

If you specify an absolute path for a file on a remote host, you should check that this is the file you intended to use. There is no guarantee that the file that appears at the specified location on the execution host is the same as or has the same contents as one that appears at the specified location on the submission host; nor is there a requirement that a file of this name be accessible on the submission host.

Specifying file names with absolute paths is useful if you have large files that you can copy to the desired location before submitting the job. For example, you might want to run several Glide docking jobs with the same grid. You could copy the grid files to the desired location, then specify them with an absolute path. Copying large files prior to job submission can help reduce network traffic, especially on clusters.

- **Files specified using relative paths with `./`**—These files are copied from their locations on the submission host to the job directory on the execution host. For example, the file `./prime/1mcp.mae` is copied to `jobdir/1mcp.mae`.
- **Files in the submission directory or its subdirectories**—These files are copied with their relative path to the job directory on the execution host. Any subdirectories are created in the job directory. For example, `1mcp.mae` is copied to `jobdir/1mcp.mae`, and `prime/1mcp.mae` is copied to `jobdir/prime/1mcp.mae`.

Output files are copied back from the job directory to the submission directory.

2.9 Incorporation of Job Output

When Maestro project entries are used as input for jobs, the structure and property output can be incorporated into the project. The output files are always copied back to the submission directory: the incorporation options affect what is done in the Maestro project. You can specify how you want the results to be incorporated in the Job Settings dialog box (see [Section 2.2 on page 5](#)). To open the Job Settings dialog box, click the button in the panel used to set up the job. The incorporation options and their effects are described below.

- **Append new entries as a new group**—New entries are added to the end of the entry list in the project as an entry group. The group name is the name of the file from which the entries originated, minus the extension. If there is only one entry, it is appended as an individual entry, not as a group. This is the default for jobs submitted from Maestro. Corresponds to the command-line option `-DISP append`.
- **Append new entries in place**—Each new entry is added to the project immediately below its source entry, i.e. the entry that was used as input for the job. If there is no source entry, the new entry is added to the end of the entry list in the project as an individual entry. Corresponds to the command-line option `-DISP appendinginplace`.
- **Append new entries individually**—New entries are added to the end of the entry list in the project as individual entries. Corresponds to the command-line option `-DISP appendungrouped`.
- **Replace existing entries**—Incoming entries replace the existing entries from which they originated. Matching of the incoming and original entries is done by entry ID. Corresponds to the command-line option `-DISP replace`.
- **Replace Workspace**—The Workspace is replaced with the first structure from the output file, but no project entries are created or modified. Corresponds to the command-line option `-DISP workspace`.
- **Do not incorporate**—Structure and property data are not incorporated into the project. You can import the results later using the Import panel. See [Chapter 3](#) of the *Maestro User Manual* for more information on importing structures. Corresponds to the command-line option `-DISP ignore`.

You can also specify a project and incorporation mode when you run jobs from the command line, by using the `-PROJ` and `-DISP` options—see [Table 2.1 on page 10](#). The default for `-DISP` is `ignore` for jobs run from the command line, since the default for `-PROJ` is not to specify a project. The `ignore` option for `-DISP` is the same as the Do not incorporate option in Maestro.

By default, incorporation of job output into the project takes place only from monitoring mode in Maestro. If you are monitoring a job launched from the current project, and that job finishes during the monitoring session, incorporation is immediate. If you are not in monitoring mode when the job finishes, you must monitor the job in the Monitor panel to incorporate the results. This applies also to jobs run from the command line and associated with a project: you must monitor the job in the Monitor panel for incorporation to take place.

The default incorporation behavior can be changed in the Jobs tab of the Preferences panel (see [Section 14.13](#) of the *Maestro User Manual*). You can choose to have jobs incorporate automatically, in which case jobs launched from the current project are incorporated as soon as they finish and are ready for incorporation. This means that you will not be able to use Maestro until the job has finished incorporating. You can also choose to receive a prompt for incorporation when a job is ready. If you chose not to incorporate the job at that time, you must monitor the job to incorporate it, or use the Incorporate All button in the Monitor panel.

When a job is incorporated, by default the first structure in the job output is displayed in the Workspace. For entries that are incorporated as an entry group, the entry list is scrolled to the first of these entries. You can change this behavior to minimize disruption to your work by setting a preference in the Jobs tab of the Preferences panel.

Incorporation can be undone. Undoing incorporation removes the new entries from the project and restores the Workspace to the state it was in just before monitoring mode was entered.

Computational programs normally propagate entry names (as well as other properties) from their input into their output structures. Where the relationship is one-to-many, as in a conformational search, consecutive structures in the output file have identical entry names. If the job is run with the “append” incorporation mode, the names are identical to that of the input entry. If the entries do not have a name, they are assigned a name that is constructed from the job name, with a “dot suffix.” Thus, the output of a job named `ligand1` would be named `ligand1.1`, `ligand1.2`, and so on. Note that the entry name is no longer used as a unique identifier in the project: the entry ID is used instead.

Managing Jobs

The Job Control facility provides tools for monitoring and controlling the jobs that it runs. These tools have both a graphical user interface in the Maestro Monitor panel and a command-line interface in the `jobcontrol` command. The interface gets information from a job database that is set up for each user, and uses this information to perform various tasks, like providing job status and killing jobs.

You can monitor jobs that are run locally, run remotely, or run in a batch queue. Jobs submitted to a grid computing manager cannot be monitored, because Job Control does not have information from the grid manager. A script has been provided to obtain information on jobs running on a United Devices grid. Certain kinds of subjobs cannot be monitored: if they are launched from a master job that runs on a remote host, and the submission directory on the remote host is not mounted on the local host, the information is unavailable.

The first section in this chapter describes the job database. Subsequent sections describe the Monitor panel and the `jobcontrol` command. Use of the `jobcontrol` command is not supported on Windows hosts.

3.1 The Job Database

Information about each job is kept in the user's job database. The database contains a record for each job. You can determine which jobs are in the database by using the Maestro Monitor panel or the command-line `jobcontrol` utility.

By default, this database is kept in the directory `$HOME/.schrodinger/.jobdb2` on Linux and `%LOCALAPPDATA%\Schrodinger\.jobdb2` on Windows. If you want to change the location, you can set the environment variable `SCHRODINGER_JOBDB2` to the desired location. You should set this environment variable globally, in your `.bashrc` or `.cshrc` file on Linux, and in your system properties on Windows, so that you are always using the same job database. If you set this environment variable locally, Job Control will not be able to locate your jobs and they might not finish. (See [Appendix A](#) of the *Installation Guide* for information on setting environment variables on Windows.)

You must ensure that this database can be read and written by a job running on any host, either by making the directory directly available on the host, or by ensuring that the host has access to a host on which it is available by passwordless `ssh`. If a job does not have access to the database, it will not be updated.

Note: The job database format changed in Suite 2011, and is not compatible with previous releases. You must *not* set the job database directory to a location used for releases prior to Suite 2011.

3.1.1 The Job Record

Each job has a job record in the job database. The job record is a list of fields, one on each line, each consisting of a field name and its value. Many of these fields contain information that is only useful to the job control system, but a number may also be useful to users. Some of the latter are listed in [Table 3.1](#)

Table 3.1. Fields used in the job record.

Field Name	Meaning
BackendPid	The PID for the program carrying out the calculation
Command	The command used to start up the actual calculation
Dir	The submission and output directory (on Host)
Envs	Environment variable settings for the job
Errors	Error messages from the job control system
ExitCode	Exit code for the job. Usually reflects the exit code for the computational program. See Table 3.2 .
ExitStatus	The reason the calculation stopped
Host	The machine from which job was launched
InputFiles	Files copied from the submission directory to the job directory at startup
JobDir	The directory in which the job is run
JobHost	The machine on which the job is run
JobId	The job's JobId
JobPid	The PID for the job's <code>jmonitor</code> process, or the job's <code>jlaunch</code> process before <code>jmonitor</code> has started.
JobUser	The user name under which the job is run
LaunchTime	The time at which the job was submitted
LogFiles	Monitoring files that grow throughout the job
MonitorFiles	Files copied from the job directory to the output directory during monitoring
MonitorInterval	The interval in seconds between monitoring updates (0 if off)

Table 3.1. Fields used in the job record. (Continued)

Field Name	Meaning
Name	The job name
OutputFiles	Files copied from the job directory to the output directory at exit
ParentJobId	The JobId of the parent job, if this is a subjob
Processors	The number of processors used for a parallel or distributed job
Program	The program name
Project	The project name
StartTime	The time at which the calculation started running
Status	The current job status
StatusTime	The time at which Status was last updated
StopTime	The time at which the calculation stopped
StructureMonitorFile	The name of the monitoring file holding the molecular structure
StructureOutputFile	The name of the file holding the final molecular structure
Summary	Description of current job status
User	The user name under which job was launched
Warnings	Warning messages from the job or job control

To list the complete database record for a job, enter the command:

```
$SCHRODINGER/jobcontrol -dump jobid
```

Common values of ExitCode for the job are given in [Table 3.2](#). On Linux, exit codes with a value greater than 128 are calculated by adding 128 to the signal number generated by the executable; thus exit code 137 corresponds to signal 9 (the kill signal).

3.1.2 Job Status

Every job proceeds through a series of well-defined stages (see [Section 2.1 on page 3](#)). The current stage of a job is displayed in the Status column of the Monitor panel. The Status of a completed job indicates the conditions under which it stopped. The job status descriptors that can appear in the Monitor panel are listed in [Table 3.3](#). When program execution finishes, an exit status is returned, and is displayed with the completed status and the incorporated status in the Monitor panel. The exit status descriptors are listed in [Table 3.4](#).

Table 3.2. Common job exit codes

Exit code	Meaning
0	Success
1	Failure, unspecified
15	Retriable licensing failure, such as inability to contact server
16	All licenses are in use
17	Fatal licensing failure, such as license does not exist
136	Floating-point error
137	Kill signal received
139	Segmentation fault
143	Termination signal received

Table 3.3. Job status descriptors in the Monitor panel.

Status	Meaning
launched	The job was submitted and assigned a JobId.
submitted	The job is in a batch queue, waiting to be scheduled.
started	The environment for the job is being set up.
running	The program is running.
paused	The program has temporarily been suspended.
exited	The program has stopped and the job is being cleaned up.
completed	The job has been cleaned up. This descriptor is followed by a colon and the exit status (see Table 3.4).
incorporated	The job results have been incorporated into a Maestro project.
stranded	Job Control could not retrieve results or clean up the job.

Table 3.4. Exit status descriptors in the Monitor panel.

Status	Meaning
finished	The program finished successfully.
stopped	The program was stopped at an appropriate point at the user's request.
killed	The job was killed by someone (not necessarily the user).
died	The program failed during execution.
fizzled	The job failed before the program could be run.

3.2 Managing Jobs From Maestro

You can manage jobs from the Monitor panel or from the command line. The common tasks can be performed from the Monitor panel, but the full range of job control tasks is only available from the command line. A web interface is also available for monitoring job progress, which can be started with `jobcontrol -g`.

To open the Monitor panel, do one of the following:

- Choose Applications → Monitor Jobs
- Choose Tasks → Monitor Jobs
- Click the Jobs button on the status bar in the main window
- Click the Monitor button in the job status window that opens from the Job toolbar.
- Double-click a row in the job status window.

The Monitor panel opens automatically when a job starts if you have the preference set under Jobs – Monitoring in the Preferences panel.

The Monitor panel provides an interface to the job database. The panel has a jobs table and two tabs: Details, and File. Below the table is a row of action buttons. The table and the tabs are described in the following subsections.

When a job is started, by default Maestro immediately goes into monitoring mode. If the preference is set, the Monitor panel opens, and the log file for the job is displayed in the text area of the File tab. The display is updated as the log file changes. For some jobs, such as MacroModel and Jaguar geometry optimizations, the Workspace is updated with each new geometry as it is generated. You can choose whether to enter monitoring mode, whether to display the Monitor panel, and set other job-related values in the Preferences panel—see [Section 14.13](#) of the *Maestro User Manual*.

Panels that do not have a Job toolbar, such as the Prime – Structure Prediction and Phase – Develop Pharmacophore Model panels, have a button that indicates whether a job is running.

The icon rotates when the job is running, and stops rotating when the job has finished. For these panels, the Monitor panel can be opened by clicking the button.

Most Maestro operations stop the monitoring of a job. You can resume or begin monitoring a job at any time by selecting it from the list of jobs in the Monitor panel.

3.2.1 The Jobs Table

The Jobs table lists jobs started by the current user. The job information is taken from the job database for the user and is updated periodically. The table is presented as a tree, in which the subjobs of a job can be hidden or displayed by clicking on the “turner” for the job (and likewise if subjobs themselves have subjobs). You can select multiple rows in the table with the usual shift-click and control-click actions. If you select multiple rows, the Details and File tabs and the Monitor button are unavailable. You can apply an action to the selected jobs with the other action buttons. You can sort the table rows by clicking the heading of the column whose values you want to sort by. Subjobs are sorted in chronological order by default.

The table cells have tool tips that display the full content of the cell, except for the Status column, where an informative message describing the meaning of the status is displayed, and the Host column, where the host entry name from the hosts file is displayed.

All jobs started by a user can be monitored in the Monitor panel, regardless of whether they were initiated from Maestro or from the command line. The class of jobs displayed can be selected from the Show option menu: all jobs, active jobs, or jobs from the current project.

To monitor a job, select the table row and click Monitor, or double-click the table row. The table row is highlighted in blue, and the File tab is placed on top with the log file displayed.

The Monitor panel allows jobs to be paused, resumed, or terminated using the Pause, Resume, Stop, and Kill buttons. Clicking Kill terminates the selected computational job immediately. The Stop button only affects MacroModel, Jaguar, and Desmond jobs, and is ignored by other jobs. When a job is stopped, the program saves results from the current stage of execution and cleans up before exiting. Pause and Resume can be used for all jobs.

If a job fails, a postmortem archive of the job can be created for sending to technical support, by using the Postmortem button. The Create Postmortem dialog box opens, in which you can choose whether to include the structures in the archive (Include structures), whether to automatically modify path names in the files so they are unrecognizable (Automatically obfuscate path names), and create the archive. For each selected job, an archive is created and written to the current working directory on Linux or the Desktop on Windows. The archive is created with the `postmortem` utility, which is described in [Appendix A](#). If you are concerned about sending confidential information, you can use this utility directly to control the content of the archive. An information dialog box provides the names of the archive files, including the path.

You can clean up the job database by clicking the Clean Up button. Jobs that are completed and incorporated, or finished if they are not incorporatable, are removed from the database.

If you want to update the status of all running jobs, click Refresh.

To incorporate all completed jobs for the current project, click Incorporate All. Jobs that were submitted with incorporation method set to Workspace are incorporated, but no structure is displayed in the Workspace.

The frequency with which information on the currently monitored job is updated can be set with a preference. You can open the Preferences panel at the Jobs - Monitoring tab by clicking Preferences.

3.2.2 The Details Tab

The Details tab gives details of the job that is selected in the Jobs tab. The job information is repeated in the Job Summary text area. The Files table lists all files associated with the job except for the structure files. When you select a file in the table, the file is displayed in the File tab. If you double-click a file, the File tab is displayed with the selected file. By default, the log file is selected (or the first log file if there is more than one). Gray rows in the table mark files that are not available for viewing. The job record for the job is the last file in the list. This file shows details of the progress of the job and is useful for diagnosing errors.

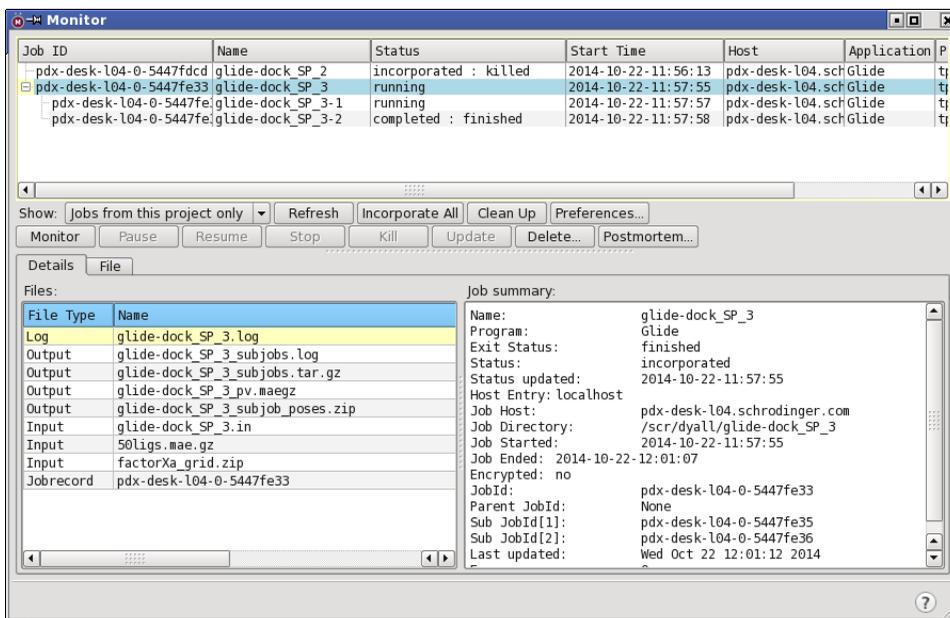


Figure 3.1. The Monitor panel, showing the Details tab.

3.2.3 The File Tab

The File tab displays the file that is selected in the Details tab. By default, this file is the log file for the job (or the first log file if there is more than one). The display is updated at regular intervals. When new text is added, the display scrolls to show any new text if the display is already at the end of the file. If the display is not at the end of the file when text is added, the display remains at the same location in the file, and you must scroll down to the end of the file to see the new text.

3.3 Managing Jobs From the Command Line

The job control utility allows you to perform a number of job control tasks from the command line. The syntax for job control utility commands is:

```
$SCHRODINGER/jobcontrol action [options] query
```

where *action* is the command for the action you want to perform, *options* qualifies the scope of the command, and *query* defines the scope of the action performed by the command. For information on the actions and options, run `jobcontrol -h`.

The optional *query* consists of one or more JobIds, job names, status codes or queries, or the keywords `all`, `active`, or `lastN`. The default *query* is all active jobs (jobs that are not finished), and is equivalent to using the keyword `active`. The keyword `all` means all jobs in your job database. The keyword `lastN` means the *N* most recently submitted jobs; `last` is equivalent to `last1`. The JobId is a unique identifier consisting of the name of the submission host, a sequence number, and a hexadecimal timestamp, e.g., `mirabelle-0-a1b2c3d4`. The job record fields and their meanings are listed in [Table 3.1 on page 22](#).

The following examples illustrate different values of *query*.

- To list all active jobs, showing their JobIds, job names, current status and the machine on which each is running, enter:

```
$SCHRODINGER/jobcontrol -list
```

- To list all the jobs in your job database that finished successfully, enter:

```
$SCHRODINGER/jobcontrol -list finished
```

- To list just the job whose JobId is `serv01-0-a1b2c3d4`, enter:

```
$SCHRODINGER/jobcontrol -list serv01-0-a1b2c3d4
```

- To list all jobs in your database with the job name `myjob`, enter:

```
$SCHRODINGER/jobcontrol -list myjob
```

- To list all jobs in your database, enter:

```
$SCHRODINGER/jobcontrol -list all
```

- To list the 5 jobs most recently submitted, enter:

```
$SCHRODINGER/jobcontrol -list last5
```

You can also use the wildcard character ‘?’ to match a single unspecified character, or ‘*’ to match zero or more unspecified characters. If you use either of these characters, you must protect them to ensure that they are interpreted by the `jobcontrol` script and not the UNIX shell. For example, you could enter either of the following commands to list all jobs whose job names start with `docklig`

```
$SCHRODINGER/jobcontrol -list docklig\  
$SCHRODINGER/jobcontrol -list 'docklig*'
```

3.3.1 General Job Control Queries

More general queries than those given above are also possible. Formally, a query consists of one or more *conditions*, optionally separated by the Boolean operators AND, OR or NOT. If the operators are omitted, OR is assumed. A condition takes the form `<field><op><value>`, where `<field>` is one of the field names in the job record and `<op>` is one of the following:

```
= equals  
!= is not equal to  
=~ matches  
!~ does not match
```

“Equals” means an exact match; “matches” means that `<field>` matches `<value>`, treated as a regular expression.

The field names are listed in [Table 3.1 on page 22](#). The fields most likely to be useful for queries are Name, Program, Host, Dir, JobHost, JobDir, Project, Status, and ExitStatus. The case of the field names is ignored, but the case of the `<value>` is significant. So, `Program=Jaguar` is the same as `program=Jaguar`, but `program=jaguar` would fail. Parentheses can be used to group conditions, but these must also be protected if used on the Unix command line.

For example, to list all QSite jobs in your database that either died or were killed, enter the command:

```
$SCHRODINGER/jobcontrol -list program=QSite AND  
  \ ( died OR killed )
```

3.3.2 Recovering Stranded Jobs

Job Control periodically tries to re-establish contact with stranded jobs. It is also possible to re-establish contact with the job manually, with the command:

```
$SCHRODINGER/jobcontrol -recover jobid
```

If this command fails, however, the job is cleaned up and the results are no longer available.

If you cannot recover a job you can remove the job record with the command:

```
$SCHRODINGER/jobcontrol -delete -force jobid
```

This command should only be used as a last resort, because it removes all job control information for the given job. However, if the job is actually alive, its job record may eventually reappear after using this command.

3.3.3 Using jserver

The `jserver` process on the submission host sets up and manages the connection to remote hosts. Normally it is started automatically. If it has stopped, you can restart it on a UNIX system with the following command:

```
$SCHRODINGER/utilities/jserver [-proxy] [action] [options]
```

For information on the actions and options, run `jserver -h`. The options only apply to newly started `jserver` processes.

In addition to starting a new `jserver` process, the `jserver` command can be used to interact with a currently running `jserver` process. If a new `jserver` process is started without the `-force` action, it checks whether another `jserver` process is already running. If so, the new process contacts the running process to determine which process has the latest version. If the new process has a later version, it kills the old `jserver` process and replaces it; otherwise, the new process exits.

A `jserver` process started with the `-proxy` option is referred to as a `jproxy` process. These processes are used for batch jobs to mediate communications between the job and the `jserver` process running on the launch host. A user can have at most one `jserver` process and one `jproxy` process running on a given host.

If `jserver` stops running, the jobs it manages are not incorporated until `jserver` restarts. Then the remote job can copy files back to the submission host. To check if `jserver` is running, use the `-info` option. You can also kill `jserver` with the `-kill` option, and you can kill `jserver` and start a new `jserver` with the `-force` option.

On Windows, `jserver` is automatically restarted when you start your computer.

The `jserver` process writes to a log file named `jserver.hostname.log` in the job database directory. You can limit the size of this file by setting the environment variable `SCHRODINGER_JSERVER_LOGSIZE`—see [Appendix B](#) for details.

3.3.4 Purging the Job Database

If the job database is not purged, job records can accumulate and the job database directory can become quite large over time. The `jobcontrol` utility can be used to purge the job database of records for completed jobs. For example, the following command purges the entire database:

```
$SCHRODINGER/jobcontrol -delete all
```

This command only affects completed jobs: running jobs cannot be deleted unless `-force` is supplied after `-delete`. However, you should be careful not to delete completed jobs whose output you still intend to incorporate into a Maestro project.

The job database is periodically checked for jobs that have finished by a process called `jnanny` (see [Section 3.3.5 on page 31](#)). If the job is older than a threshold time, its record is deleted from the database. You can set this threshold time with `jnanny` or with the environment variable `SCHRODINGER_JOBDB_CLEANUP`. The default time is 1 week; the minimum allowed time is 1 second, but a realistic minimum is 10 minutes. The default unit is seconds, but you can specify a time in minutes, hours, or days by appending `m`, `h`, or `d` to the value, for example, `7d`, `168h`, or `5m`.

3.3.5 Performing Actions on Jobs with `jnanny`

The `jnanny` script scans the job database and performs tasks defined by plugin modules. These tasks include checking for jobs that are stranded or appear to be stuck in a particular state, and then initiating a recovery process and updating the job database. This script is run periodically by `jserver`. The interval at which an action is performed can be defined by setting an environment variable or using the `-set` option—see [Table 3.5](#). The default unit is seconds, but you can specify a time in minutes, hours, or days by appending `m`, `h`, or `d` to the value.

Since `jnanny` is run periodically, it should not in general be necessary to run it manually. However, you can do so if the need arises, with the following command:

```
$SCHRODINGER/utilities/jnanny [options] [plugins]
```

For information on the options and plugins, run `jnanny -h`. If no plugins are specified, all plugins are called. You can abbreviate a plugin specification to the shortest unique string.

Table 3.5. Settings and environment variables for jnanny actions.

Parameter	Description
<i>CheckStranded module settings</i>	
stranded_threshold	An active job whose job record has not been updated for this length of time is checked to verify that it is still running. Default: 10m. Minimum: 1m. 0 disables checking. Environment variable: SCHRODINGER_STRANDED_THRESHOLD
stranded_interval	Interval at which jnanny checks for stranded jobs. Default: 10m. Minimum: 20s Environment variable: SCHRODINGER_STRANDED_INTERVAL
death_threshold	An active job that is unreachable for this length of time is marked dead. Default: 2h. Minimum: 10m. 0 means infinite time. Environment variable: SCHRODINGER_DEATH_THRESHOLD
<i>CleanupDatabase module settings</i>	
jobdb_cleanup	A completed job that is not part of an active distributed job is deleted from the job database after this length of time. Default: 7d. 0 turns off subsequent cleanup after cleaning up. Environment variable: SCHRODINGER_JOBDB_CLEANUP
forced_cleanup	All jobs that have been in the database longer than this length of time are deleted from the job database, regardless of status. Default: 2 months. 0 turns off subsequent cleanup. Environment variable: SCHRODINGER_FORCED_CLEANUP
jobdb_maxsize	If there are more than this many jobs in the job database, completed jobs that are not part of an active distributed job are cleaned up to lower the number of jobs below this limit. Default: 2000. Minimum: 0. Environment variable: SCHRODINGER_JOBDB_MAXSIZE
logfile_cleanup	Remove a rotated jserver or jproxy log file that is older than this value. Default: 30d. Minimum: 1m. Environment variable: SCHRODINGER_LOGFILE_CLEANUP
<i>CheckSubmitted module settings</i>	
submitted_interval	Interval at which jnanny checks for stranded jobs that are submitted. Default: 10m. Minimum: 1m. Environment variable: SCHRODINGER_SUBMITTED_INTERVAL
submitted_threshold	A submitted job that has not started running after this length of time is checked to make sure it is still on the queue. Default: 5m; Minimum: 1m Environment variable: SCHRODINGER_SUBMITTED_THRESHOLD

The postmortem Utility

The `postmortem` utility archives information that is useful for understanding why jobs do not run as expected. It creates a zip archive containing the current environment, file system information, a list of installed Schrodinger software packages, the `schrodinger.hosts` file, the queue support scripts, the license file, and information such as input and output for specified jobs. When contacting technical support regarding a job, it is highly recommended to run this utility, and upload the archive when you fill in the support request form on the web site (<https://www.schrodinger.com/supportcenter/>). It is also highly recommended that you generate this archive using the Diagnostics panel, which runs the `postmortem` utility and creates the archive with other information as well.

Note: You should ensure that no sensitive information is added to the archive, by using the options described below, and by inspecting the contents of the archive.

The syntax of the `postmortem` utility is as follows:

```
$SCHRODINGER/utilities/postmortem [options] [jobids]
```

The options are described in [Table A.1](#). Options may be truncated to any unique abbreviation.

Table A.1. Options for the postmortem command.

Option	Description
<code>-autoreplace</code>	Automatically replace strings that may contain sensitive information.
<code>-exclude <i>suff</i></code>	Do not include files with the given suffix in the archive.
<code>-help</code>	Show the usage message.
<code>-[no] alljobs</code>	Include information on all jobs in the archive. Default: include information on specified jobs only.
<code>-[no] parents</code>	Include [exclude] parent jobs of the specified jobs. Default: include.
<code>-[no] subjobs</code>	Include [exclude] subjobs of the specified jobs. Default: include
<code>-[no] jobfiles</code>	Include [exclude] input, output and log files. Default: include.
<code>-[no] datafiles</code>	Include [exclude] data files. Default: exclude.
<code>-output <i>name</i></code>	The base name of the archive: the archive is named <code><i>name</i>.zip</code> . The default is <code><i>username-host-schrodinger</i>.zip</code> if no job IDs are given, and <code><i>jobid1-<i>archive</i>.zip</i></code> if job IDs are given, with <code><i>jobid1</i></code> the first job ID.

Table A.1. Options for the postmortem command.

Option	Description
-pref <i>name</i> [=value]	Set preferences for the default behavior of postmortem. The allowed names and values, with the defaults, are given in Table A.2 . If the value is omitted, display the current value of the preference.
-prefs	List the currently defined preferences.
-quiet	Do not print anything on standard output.
-replace <i>string</i>	Replace the given string in the archived files.
-reset [<i>name</i>]	Reset the specified preference to its default value. If name is not given, reset all preferences to their default values.
-size	Calculate the total size of the job files and exit; do not create an archive.
-[no]structures	Include [exclude] structure files. Default: exclude.
-verbose	List the names of the archived files on standard output.

If job IDs are specified, the job records for those jobs are archived, along with any related files (batch scripts, qlog files, and so on) from the job database. Parent and subjobs of the specified jobs are included as well, unless the option `-noparents` or `-nosubjobs` is given. Finally, unless `-nojobfiles` is given, any log files and non-structural input or output files listed in the job record are saved as well. If you want to include structure files in the output, use the `-struct` option.

The archive can be quite large if all job files for a large job are included. You might want to use the `-size` option to check the total size of the files to be archived, first.

To find the job IDs for your jobs, use the `jobcontrol` command, or look in the Monitor panel.

For example the following command lists all of your completed jobs:

```
$SCHRODINGER/jobcontrol -list completed
```

If file and directory names contain sensitive information that you don't want to reveal, you can use the `-autoreplace` option to have the program replace them in the archived files. You can specify particular string replacements using the `-replace` option, as well. A list of all string replacements that were done is written to a file called `archive.names`.

You can change the default behavior of `postmortem` by setting preferences. This allows you to run the `postmortem` command without any arguments and archive your preferred choice of information. These preferences are stored in your Schrödinger user resources directory, and are managed using `postmortem` command options. The preference names and values are given in [Table A.2](#). Unlike options, preference names must not be abbreviated.

Table A.2. Preferences for the postmortem utility, with the corresponding option

Preference	Option	Values	Default
location	-location	String defining directory	Current directory
jobfiles	-[no]jobfiles	1 0	1 (-jobfiles)
alljobs	-[no]alljobs	1 0	0 (-noalljobs)
structures	-[no]structures	1 0	0 (-nostructures)
datafiles	-[no]datafiles	1 0	0 (-nodatafiles)
parents	-[no]parents	1 0	1 (-parents)
subjobs	-[no]subjobs	1 0	1 (-subjobs)
verbose	-verbose	1 0	1 (-verbose)

- To set preferences, use the `-pref` option:

```
postmortem -pref structures=1
postmortem -pref parents=0
```
- To see what preferences are currently in force, use the `-prefs` option,

```
postmortem -prefs
```
- To see the value for a particular preference, use `-pref` without assigning a value:

```
postmortem -pref verbose
```
- To remove a preference, use the `-reset` option, for example

```
postmortem -reset verbose
```
- To remove all preferences, use `-reset` without specifying a preference name:

```
postmortem -reset
```


Environment Variables

There are a number of environment variables that can be used to specify the location of resources, manage the job database, and control various aspects of job execution. These environment variables are given in [Table B.1](#). Some of these environment variables provide alternate means of specifying the resource; others are the sole means of specifying a resource or of overriding a default value. For information on setting environment variables, see [Appendix A](#) of the *Installation Guide*.

Table B.1. Environment variables for resource specification and job control.

Variable	Purpose
SCHROD_LICENSE_FILE LM_LICENSE_FILE	Specify the license file or a list of license files to use if the license file is not installed in \$SCHRODINGER. The list must be separated by colons on Unix and semicolons on Windows. The SCHROD_LICENSE_FILE definition supersedes any other definition. Both of these environment variables are FlexLM environment variables. Use SCHROD_LICENSE_FILE when LM_LICENSE_FILE does not point to a license file containing Schrödinger licenses.
SCHRODINGER	Specify the installation directory.
SCHRODINGER_AUTO_SAVE	If set to 1, force the contents of the job directory to be archived and copied back to the submission host only for those jobs that have died (whose exit status is “died”). It is also possible to force such copying for all jobs (see SCHRODINGER_SAVE_JOBDIR).
SCHRODINGER_DEATH_THRESHOLD	Time threshold after which an active but unreachable job is considered dead.
SCHRODINGER_HOSTS	Specify the hosts file. See Section 2.3.5 on page 15 .
SCHRODINGER_JOB_DEBUG	Control debugging output. Allowed values are: <ul style="list-style-type: none"> 0 No debugging output 1 Standard debugging output (same as entering -DEBUG from the command line) 2 Detailed debugging output (same as entering -DDEBUG from the command line)

Table B.1. Environment variables for resource specification and job control. (Continued)

Variable	Purpose
SCHRODINGER_JOB_DISPOSITION	Specify how to incorporate job results into the project (same as entering <code>-DISP</code> from the command line) Available values are: append—add entries to project in an entry group appendungrouped—add entries to project singly replace—replace project entries with new entries ignore—do not incorporate entries Default: ignore See Section 2.9 on page 19 .
SCHRODINGER_JOBDB2	Specify the full path name for the job database. Must not be set to a location that contains a pre-2011 job database. Note that SCHRODINGER_JOBDB is no longer supported, and will be ignored. Default: <code>~/schrodinger/.jobdb2</code> See Section 3.1 on page 21 .
SCHRODINGER_JOBDB_CLEANUP	Specify how long the job record of a completed job is kept in the job database before being automatically deleted. For example, one day can be expressed as <code>86400</code> , <code>1440m</code> , <code>24h</code> , or <code>1d</code> . A value of <code>0</code> turns off automatic cleanup. The value is overridden by any <code>jnanny</code> setting, so you should use <code>jnanny</code> instead if possible. Default: 7 days. See Section 3.3.4 and Section 3.3.5 on page 31 .
SCHRODINGER_JOBDB_MAXSIZE	Specify the maximum number of job records to be kept in the job database. If this number is exceeded, job records of completed jobs are purged, starting with the oldest, at the next cleanup time. Default: 2000 records.
SCHRODINGER_JMONITOR_PORT	Specify the port or range of ports that <code>jmonitor</code> may use. Ports can be specified as comma or colon-separated lists without spaces. Ranges can be specified with a dash, for example, <code>5987:5989-5992:5994</code> . Set this variable if there is a firewall that allows connections to ports within a given range only. The number of ports should be twice the number of <code>jmonitor</code> processes (jobs).
SCHRODINGER_JPROXY	Regulate the use of <code>jproxy</code> for the job. If set to <code>1</code> , <code>jproxy</code> is used; if set to <code>0</code> , it is not. By default (if not set), <code>jproxy</code> is used for batch jobs.

Table B.1. Environment variables for resource specification and job control. (Continued)

Variable	Purpose
SCHRODINGER_JPROXY_PORT	Specify the port or range of ports that jproxy may use. Overrides any setting of proxyport in the hosts file. The syntax is the same as for SCHRODINGER_JSERVER_PORT. If this environment variable is not set, the values set in SCHRODINGER_JSERVER_PORT apply to jproxy as well. Therefore, it is generally not advisable to specify only one port in SCHRODINGER_JSERVER_PORT, as this will make it impossible for jserver and jproxy to run on the same machine, which may be necessary.
SCHRODINGER_JSERVER	Regulate the use of jserver for the job. If set to 1, jserver is used; if set to 0, it is not. By default (if not set), jserver is used for remote and batch jobs.
SCHRODINGER_JSERVER_LOGSIZE	Set the size threshold for the jserver log file in bytes. Units of kilobytes and megabytes can be used by appending K or M to the value. The five most recent log files are kept, with extensions .1 through .5. When the file size exceeds this threshold, the files are rotated, the .log.5 file is discarded if it exists, and a new log file is started.
SCHRODINGER_JSERVER_PORT	Specify the port or range of ports that jserver may use. Ports can be specified as comma or colon-separated lists without spaces. Ranges can be specified with a dash, for example, 5987:5989-5992:5994. Setting this variable is useful if there is a firewall that allows connections to ports within a given range only.
SCHRODINGER_LICENSE_DEBUG	Turns on debugging related to license usage. Any non-zero integer value turns on debugging.
SCHRODINGER_LICENSE_RETRY	Specify the maximum time to keep trying to obtain a license. The value can be an integer value in seconds, or an integer value with a time unit, such as 7200s, 120m, 2h. An attempt to obtain a license is made at intervals that double after each try, starting at 10s. The final attempt is made at the specified maximum time. Default: 1 minute.
SCHRODINGER_LOGFILE_CLEANUP	Clean up (remove) jserver and jproxy log files that are older than the specified age. The value can be an integer value in seconds, or an integer value with a time unit, such as 7200s, 120m, 2h. Default: 1 month.

Table B.1. Environment variables for resource specification and job control. (Continued)

Variable	Purpose
SCHRODINGER_MAX_RETRIES	Specify the number of times a failed subjob from a distributed job is rerun before exiting with a failure. Only supported for distributed Glide and workflow jobs. Default: 3.
SCHRODINGER_MONITOR_INTERVAL	Set the interval at which monitored files are copied back to the launch directory.
SCHRODINGER_NICE	Lower the priority of Schrödinger jobs. If it is unset or set to a null string, the priority is not lowered. The value <code>local</code> lowers the priority on jobs and subjobs submitted to the local host, but does not lower the priority of jobs or subjobs submitted to a remote host. Any other value lowers the priority of all jobs. By default, most jobs run on the local host from Maestro are run with reduced priority. To change this, set this variable to a null string. Equivalent to using the <code>-NICE</code> command-line option.
SCHRODINGER_NO_NFS	If set, avoid using NFS mounts for file transfers between the job directory and the submission (output) directory, if possible.
SCHRODINGER_NODEFILE	Specify the path to a file containing a list of hosts and the number of CPUs to use on each host. Each line in the file must have the format: <i>host: numberofCPUs</i> You can list a host multiple times rather than provide the number of CPUs; this may affect the execution of parallel jobs. This environment variable is used to pass an explicit list of hosts to the driver for a parallel or distributed job. It is not used by Job Control directly; as such, the meaning of the file depends on the product that uses it. For a parallel job, such as parallel Jaguar or Desmond, the names in the file must be actual host names, not entries in the hosts file. For distributed jobs, like Glide docking, the names must be entries in the hosts file, since they are used for launching subjobs. Ordinarily, the hosts to use for a distributed or parallel job would be supplied using command line arguments, as explained elsewhere in this document. This variable is provided to allow the usual mechanism to be overridden when necessary.

Table B.1. Environment variables for resource specification and job control. (Continued)

Variable	Purpose
SCHRODINGER_PROJECT	Specify the path name of the Maestro project into which the results of the job should be incorporated (same as entering <code>-PROJ</code> from the command line) See Section 2.9 on page 19.
SCHRODINGER_QLOGDIR	Specify the path to the directory where queue log files (<code>.qlog</code>) for batch jobs are to be written. The default is the job database directory; however, if that directory is not mounted on the queue host, an alternative path needs to be specified, or the queuing system will not be able to write those log files.
SCHRODINGER_QUEUE_ARGS	Specify queue arguments for job submission to a queuing system. These arguments are also used when subjobs are submitted by a master job. Queue arguments specified with this environment variable are appended to those read from <code>schrodinger.hosts</code> .
SCHRODINGER_RETAIN_JOBDIR	If set to 1, when the job finishes and the output files are copied back, the job directory is not removed. By default, it is removed if it did not exist already.
SCHRODINGER_SAVE_JOBDIR	If set to 1, archive the contents of the job directory and copy it back to the submission host when the job finishes. The archive appears in the output directory as <code>jobid-jobdir.zip</code> . It is also possible to force such copying only for jobs that have died (see <code>SCHRODINGER_AUTO_SAVE</code>). Equivalent to using the <code>-SAVE</code> command-line option.
SCHRODINGER_SECURE_TRANSFERS	Enforce or cancel data transfer via <code>ssh</code> tunnels. If set to 1, secure transfers are used regardless of any secure zone settings in the hosts file. If set to 0, secure transfers are not used at all. If unset, secure transfers are used for transfers outside the secure zones defined in the hosts file.
SCHRODINGER_STRANDED_INTERVAL	Time interval for checking for stranded jobs by <code>jnanny</code> . Default: 10 minutes.
SCHRODINGER_STRANDED_THRESHOLD	Check whether an active job is still running if its job record has not been updated for this length of time. Default: 10 minutes.
SCHRODINGER_SUBMITTED_INTERVAL	Time interval for checking on submitted and launched jobs by <code>jnanny</code> to ensure that they are progressing. Default: 10 minutes.

Table B.1. Environment variables for resource specification and job control. (Continued)

Variable	Purpose
SCHRODINGER_SUBMITTED_THRESHOLD	Check whether a submitted job is still in the queue if its job record has not been updated for this length of time. Default: 5 minutes.
SCHRODINGER_TMPDIR	Specify the full path name for the directory in which run time temporary job directories are created. Overrides the tmpdir setting specified in the schrodinger.hosts file (Section 7.1.3 of the <i>Installation Guide</i>).

Secure File Transfer to Queue Hosts

All file transfers made to and from queue hosts for jobs submitted to a batch queue are secured by default, so that sensitive information can be transmitted across insecure networks. The core of the secure transfer mechanism is an SSH tunnel (“secure channel”) established between `jserver` on the launch host and `jproxy` on the queue manager host. The SSH tunnel is established in the background, automatically, in a transparent manner, and is removed when it is no longer needed. This ensures that all data transfers between the launch host and the computing cluster are encrypted. Communication between the queue manager host and compute nodes is not secured, as the main goal of the secure file transfer mechanism is to protect information as it enters and leaves the queue host.

Simple remote jobs from one host to another, which do not involve a batch queue, are never secured by the secure transfer mechanism.

The most common use case for a secure transfer channel involves a workstation that communicates with the head node of a cluster or the Cloud via the internet. In this case, it is not desirable to transfer data in plain text, since it would be exposed over the internet. It is common for head nodes in such situations to employ firewalls that block all incoming connections except for SSH. The secure transfer mechanism is well suited for both of these concerns.

If the connection between the job submission host and the queue host is already secure (for example, part of a secure network), there is no need for the secure file transfer mechanism. In this situation secure transfers can be disabled altogether or within specified *secure zones*.

C.1 Controlling Secure File Transfer

Secure file transfers can be disabled between a specific set of hosts by defining secure zones in the hosts file. See [Section 7.1.10](#) of the *Installation Guide* for information on defining these zones.

While it is possible for a host to be in more than one secure zone, hosts never act as bridges between different zones. For example, if one zone includes `host1` and `host2` and another one includes `host2` and `host3`, then connections between `host2` and `host1` or `host3` are not secured, but connections between `host1` and `host3` are secured because they do not belong to the same secure zone. The file transfers do not take place via `host2`, so it does not act as a bridge between the two secure zones.

Secure file transfers can be turned on and off globally by setting the environment variable `SCHRODINGER_SECURE_TRANSFERS`. A value of 0 turns off secure file transfers, and a value of 1 turns on secure file transfers to all queue hosts. This setting takes precedence over all other settings, including secure zones. Secure zones are used only if this variable is undefined.

The active secure zone information is stored in the job database directory, in the same format as in the hosts file:

```
~/ .schrodinger/ .jobdb2/securezones .host
```

C.2 Altering Secure Zone Settings

Since secure zones are interpreted and used by `jserver`, which is shared between all of the jobs that you launched from a given host, you cannot change the secure zone configuration while `jserver` is active. If you do, jobs launched after such a change fail with an error message: “Conflicting secure zones information”. If you do need to change the secure zone configuration, you must kill your current `jserver`, with the command

```
$(SCHRODINGER/utilities/jserver -clean
```

and then make changes to the `securezone` entries of your hosts file before launching any subsequent jobs. If the hosts file is on a network file system, all `jserver` processes that have access to that file system must be killed.

The same applies to the `SCHRODINGER_SECURE_TRANSFERS` environment variable. If you need to change its value (including undefining it if it was defined previously), you must kill `jserver` before launching any further jobs.

C.3 Verifying Secure Transfer

To verify whether connections between a `jserver` process and one or more of `jproxy` processes are secure, use the command

```
$(SCHRODINGER/utilities/jserver -show jproxys
```

The output should be something like the following:

```
Found jserver (PID 1645) listening on port 47129
monitoring 2 jproxy processes
untrustedhost [JPROXY_OK (secure)]
trustedhost [JPROXY_OK]
```

In this example, the connection between your host and `untrustedhost` is secured, while that between your host and `trustedhost` is not.

You can also use the Diagnostics GUI to determine whether or not jobs sent to a particular entry in the hosts file is secured. In the GUI, select the desired host from the Hosts table, and use **Selected Host** to run a test job. If the `testapp` result is `OK (SECURE)`, then jobs sent to that cluster entry will use the secure transfer mechanism.

Getting Help

Schrödinger software is distributed with documentation in PDF format. If the documentation is not installed in `$SCHRODINGER/docs` on a computer that you have access to, you should install it or ask your system administrator to install it.

For help installing and setting up licenses for Schrödinger software and installing documentation, see the *Installation Guide*. For information on a particular product, see the documentation for that product. For information on using Maestro and its panels, see the Maestro online help or the *Maestro User Manual*. Other information is available on the Schrödinger [Support Center](#) and the knowledge base, <http://www.schrodinger.com/kb>.

Contacting Technical Support

If you have questions that are not answered from any of the above sources, contact Schrödinger using the information below.

Web: <http://www.schrodinger.com/supportcenter>
E-mail: help@schrodinger.com
Mail: Schrödinger, 101 SW Main Street, Suite 1300, Portland, OR 97204
Phone: +1 888 891-4701 (USA, 8am – 8pm Eastern Time)
+49 621 438-55173 (Europe, 9am – 5pm Central European Time)
Fax: +1 503 299-4532 (USA, Portland office)
FTP: <ftp://ftp.schrodinger.com>

Generally, using the web form is best because you can add machine output and upload files, if necessary. You will need to include the following information:

- All relevant user input and machine output
- Schrödinger software purchaser (company, research institution, or individual)
- Primary Schrödinger software user
- Installation, licensing, and machine information as described below.

Gathering Information for Technical Support

The instructions below describe how to gather the required machine, licensing, and installation information, and any other job-related or failure-related information, to send to technical support. Where the instructions depend on the profile used for Maestro, the profile is indicated.

For general enquiries or problems:

1. Open the Diagnostics panel.
 - **Maestro:** Help → Diagnostics
 - **Windows:** Start → All Programs → Schrodinger-2015-2 → Diagnostics
 - **Mac:** Applications → Schrodinger2015-2 → Diagnostics
 - **Command line:** `$SCHRODINGER/diagnostics`

2. When the diagnostics have run, click Technical Support.

A dialog box opens, with instructions. You can highlight and copy the name of the file.

3. Upload the file specified in the dialog box to the support web form.

If you have already submitted a support request, use the upload link in the email response from Schrödinger to upload the file. If you need to submit a new request, you can upload the file when you fill in the form.

If your job failed:

1. Open the Monitor panel, using the instructions for your profile as given below:

- **Maestro/Jaguar/Elements:** Tasks → Monitor Jobs
- **BioLuminate/MaterialsScience:** Tasks → Job Monitor

2. Select the failed job in the table, and click Postmortem.

The Postmortem panel opens.

3. If your data is not sensitive and you can send it, select Include structures and deselect Automatically obfuscate path names.

4. Click Create.

An archive file is created, and an information dialog box with the name and location of the file opens. You can highlight and copy the name of the file.

5. Upload the file specified in the dialog box to the support web form.

If you have already submitted a support request, use the upload link in the email response from Schrödinger to upload the file. If you need to submit a new request, you can upload the file when you fill in the form.

6. Copy and paste any log messages from the window used to start the interface or the job

into the web form (or an e-mail message), or attach them as a file.

- **Windows:** Right-click in the window and choose Select All, then press ENTER to copy the text.
- **Mac:** Start the Console application (Applications → Utilities), filter on the application that you used to start the job (Maestro, BioLuminate, Elements), copy the text.

If Maestro failed:

1. Open the Diagnostics panel.

- **Windows:** Start → All Programs → Schrodinger-2015-2 → Diagnostics
- **Mac:** Applications → SchrodingerSuite2015-2 → Diagnostics
- **Linux/command line:** `$SCHRODINGER/diagnostics`

2. When the diagnostics have run, click Technical Support.

A dialog box opens, with instructions. You can highlight and copy the name of the file.

3. Upload the file specified in the dialog box to the support web form.

If you have already submitted a support request, use the upload link in the email response from Schrödinger to upload the file. If you need to submit a new request, you can upload the file when you fill in the form.

4. Upload the error files to the support web form.

The files should be in the following location:

- **Windows:** `%LOCALAPPDATA%\Schrodinger\appcrash`
(Choose Start → Run and paste this location into the Open text box.)
Attach `maestro_error_pid.txt` and `maestro.exe_pid_timestamp.dmp`.
- **Mac:** `$HOME/Library/Logs/CrashReporter`
(Go → Home → Library → Logs → CrashReporter)
Attach `maestro_error_pid.txt` and `maestro_timestamp_machinename.crash`.
- **Linux:** `$HOME/.schrodinger/appcrash`
Attach `maestro_error_pid.txt` and `crash_report_timestamp_pid.txt`.

If a Maestro panel failed to open:

1. Copy the text in the dialog box that opens.
2. Paste the text into the support web form.

Glossary

entry—(1) A collection of settings in the hosts file that defines a configuration for running jobs on a given host. There can be more than one entry for a given host, with different settings.
(2) A structure and its properties in a Maestro project.

execution host—The computer that a job runs on.

hosts file—The file that contains information on available hosts that is used by Job Control. This file is usually named `schrodinger.hosts`.

job directory—The directory on the execution host to which the input files are copied from the submission host and from which the output files back to the submission host when the job is finished. This is usually a subdirectory of the scratch directory that is created by Job Control.

launch directory—The directory from which the job is started. Same as *submission directory*.

launch host—The computer that you submit a job from. Same as *submission host*.

local host—The computer that you are logged on to. Usually this is also the computer you are using for job submission.

output directory—The directory to which output files are copied at the end of the job. This is usually the same as the submission directory.

queue host—A computer that can run the queueing software. This is not necessarily the queue manager: for example, if the queuing software is installed on an NFS-mounted file system, it could be any host that has that file system.

remote host—A computer that is available to you over a network.

remote job submission—Submitting a job to a computer other than the one you are logged on to or that you are running Maestro on.

Schrödinger user resource area—location for storage of resources (files, scripts, data) used by Schrödinger software for an individual user. This location is `$HOME/.schrodinger` on Linux and `%APPDATA%\Schrodinger` on Windows.

scratch directory—The directory that is used for temporary files. The files are usually created in a subdirectory of this directory that is unique to the user and the job.

submission directory—The directory from which the job is started. This is the directory from which input files (specified without a path) are copied to the job directory, and to which output files are copied at the end of the job.

submission host—The computer that you submit a job from.

B

- batch queues
 - passing arguments for 10
 - queue log files..... 41
 - software version..... 17

C

- command options
 - information 11
 - job submission 10
 - parsing of 3, 4
 - program 11
- conventions, document..... v

D

- debugging
 - failed job output..... 37
 - license usage..... 39
- directories
 - installation 5
 - job 15
 - scratch..... 15–16
 - submission 15

E

- entries, project, incorporation of..... 19
- environment variables 37
 - jnanny action intervals..... 32
 - licensing..... 17, 37
 - order set 18
 - SCHRODINGER_HOSTS 15
 - SCHRODINGER_JOBDB_CLEANUP 31
 - SCHRODINGER_JOBDB2 21
 - SCHRODINGER_NICE 8
 - SCHRODINGER_SCRIPTS 12
 - SCHRODINGER_SECURE_TRANSFERS..... 44
 - SCHRODINGER_TMPDIR 16
 - table 37
 - TMPDIR 16
- exit codes 24

F

- field names, job record..... 22
- files, displaying 27

H

- hosts
 - listing 11
 - querying installation on 11
 - selecting for program execution 10
 - submission 3
- hosts file
 - location of..... 15
 - tmpdir setting 42
 - use by Maestro 5

I

- incorporation of job results 19
- input files 18
- installation directory used by job..... 17

J

- jlaunch script 3, 4
- jmonitor script 4
- jnanny script 31
- job database 21
 - cleanup time 32
 - creation of records in..... 3
 - location of..... 38
 - maximum age of records 32
 - maximum size 32, 38
- job directory 15, 51
 - archiving..... 41
 - retention after job completion 41
- job ID
 - assignment of..... 3
 - format 28
- job record
 - description of fields 22
 - maximum age 38
- Job Settings dialog box 7
- jobcontrol utility 28–31
 - purging job database..... 31
 - syntax 28
- jobs
 - archiving failed 33
 - cleaning up database..... 31
 - database 21
 - distributing..... 9
 - incorporation options..... 10, 19, 38
 - killing 26

life cycle	3	Q	
listing	28	queue manager	
lowering priority of.....	10, 40	passing arguments to	10
monitoring	25	secure file transfer	43
monitoring status keywords.....	23	R	
output incorporation	19	remote hosts	
pausing and resuming	26	querying installation on	11
record	22	selecting for program execution	10
remote submission	13	S	
scratch directory	15–16	Schrödinger contact information	47
stranded	30	scratch directory.....	15–16
temporary files.....	14	scratch files— <i>see</i> temporary files	
jproxy process.....	4, 30	settings, hosts file	
jserver	30	tmpdir	42
function of	4	use by Maestro	5
restarting.....	30	sockets	
L		use for file transfer	4
launch directory	15	software version	
licenses		selection for job.....	17
checkin, checkout	4	SSH tunnel, use for file transfer.....	4
error codes related to	24	Start dialog box.....	5
location of license file	17, 37	subjobs	
time period for obtaining.....	39	hosts for	15
log files, for queuing systems	41	retrying failed	40
M		specifying number of.....	9
Maestro project		unavailable for monitoring	21
assigning job to.....	10, 41	submission directory	15
incorporation of job results.....	10, 19, 38	submission host.....	3
Monitor panel		T	
Details tab.....	27	temporary files	
monitoring jobs	25, 28	location of.....	14, 15, 42
O		removal	16
options, Job Control command-line	9	writing to submission directory.....	11
output files	18	U	
P		unxutils package.....	12
port range	38, 39	user name, specifying for job.....	10
postmortem utility	33	W	
priority of jobs	10, 40	Windows, running from the command line	12
processors, specifying number for execution	8, 13		
product installation	47		

120 West 45th Street
17th Floor
New York, NY 10036

155 Gibbs St
Suite 430
Rockville, MD 20850-0353

Quatro House
Frimley Road
Camberley GU16 7ER
United Kingdom

101 SW Main Street
Suite 1300
Portland, OR 97204

Dynamostraße 13
D-68165 Mannheim
Germany

8F Pacific Century Place
1-11-1 Marunouchi
Chiyoda-ku, Tokyo 100-6208
Japan

245 First Street
Riverview II, 18th Floor
Cambridge, MA 02142

Zeppelinstraße 73
D-81669 München
Germany

No. 102, 4th Block
3rd Main Road, 3rd Stage
Sharada Colony
Basaveshwaranagar
Bangalore 560079, India

8910 University Center Lane
Suite 270
San Diego, CA 92122

Potsdamer Platz 11
D-10785 Berlin
Germany

SCHRÖDINGER