

Submit Command

Overview

submit takes a user command and executes it remotely. The objective is to allow the user to issue a command in the same manner as a locally executed command. Multiple submission mechanisms are available for job dissemination. A set of steps are executed for each job submission:

- Destination site is selected
- A wrapper script is generated for remote execution
- If needed a batch system description file is generated.
- Input files for a job are gathered and transferred to the remote site. Transferred files include wrapper script, batch description scripts.
- Progress of the remote job is monitored until completion.
- Output files from the job are returned to the dissemination point.

Command Syntax

submit command options can be determined by using the help parameter of the submit command.

```
$ submit --help
usage: submit [options]
```

options:

```
-v, --venue                Remote job destination
-i, --inputfile            Input file
-n NCPUS, --nCpus=NCPUS   Number of processors for MPI execution
-N PPN, --ppn=PPN         Number of processors/node for MPI execution
-w WALLTIME, --wallTime=WALLTIME
                           Estimated walltime hh:mm:ss or minutes
-e, --env                  Variable=value
-m, --manager              Multiprocessor job manager
-r NREDUNDANT, --redundancy=NREDUNDANT
                           Number of indential simulations to execute in
                           parallel
-M, --metrics              Report resource usage on exit
-W, --wait                 Wait for reduced job load before submission
-h, --help                 Report command usage
```

Currently available DESTINATIONS are:

```
clusterA
```

SUBMIT COMMAND

```
clusterB
```

Currently available MANAGERS are:

```
mpich-intel32
```

By specifying a suitable set of command line parameters it is possible to execute commands on configured remote systems. The simple premise is that a typical command line can be prefaced by submit and its arguments to execute the command remotely.

```
$ submit -v clusterA echo Hello world!  
Hello world!
```

In this example the echo command is executed on the venue named clusterA where jobs are executed directly on the host. Execution of the same command on a cluster using PBS would be done in a similar fashion

```
$ submit -v clusterB echo Hello world!  
(2586337) Simulation Queued Wed Oct 7 14:45:21 2009  
(2586337) Simulation Done Wed Oct 7 14:54:36 2009  
$ cat 00577296.stdout  
Hello world!
```

submit supports an extensible variety of submission mechanisms. HUBzero supported submission mechanisms are

- local - use job submission mechanisms available directly on the submit host. These include PBS and condor job submission.
- ssh - direct use of ssh with pre-generated key.
- ssh + remote batch job submission - use ssh to do batch job submission remotely, again with pre-generated key.

A site for remote submission occurs is selected in one of the following ways, listed in order of precedence:

- User specified on the command line with -v/--venue option.
- Randomly selected from remote sites associated pre-staged application.

SUBMIT COMMAND

Any files specified by the user plus internally generated scripts are packed into a tarball for delivery to the remote site. Individual files or entire directory trees may be listed as command inputs using the `-i/--inputfile` option. Additionally command arguments that exist as files or directories will be packed into the tarball. If using ssh based submission mechanisms the tarball is transferred using scp.

The job wrapper script is executed remotely either directly or as a batch job. The job is subject to all remote queuing restrictions and idiosyncrasies.

Remote batch jobs are monitored for progress. Changes in job status are reported at least every minute. Job status is reported at least every four minutes. The job status is used to detect job completion.

The same methods used to transfer input files are applied in reverse to retrieve output files. Any files and directories created or modified by the application are be retrieved. A tarball is retrieved and expanded to the home base directory. It is up to the user to avoid the overwriting of files.

In addition to the application generated output files additional files are generated in the course of remote job execution. Some of these files are for internal bookkeeping and are consumed by submit, a few files however remain in the home base directory. The remaining files include `JOBID.stdout` and `JOBID.stderr`, it is also possible that a second set of standard output/error files will exist containing the output from the batch job submission script. `JOBID` represents unique job identifier assigned by submit.