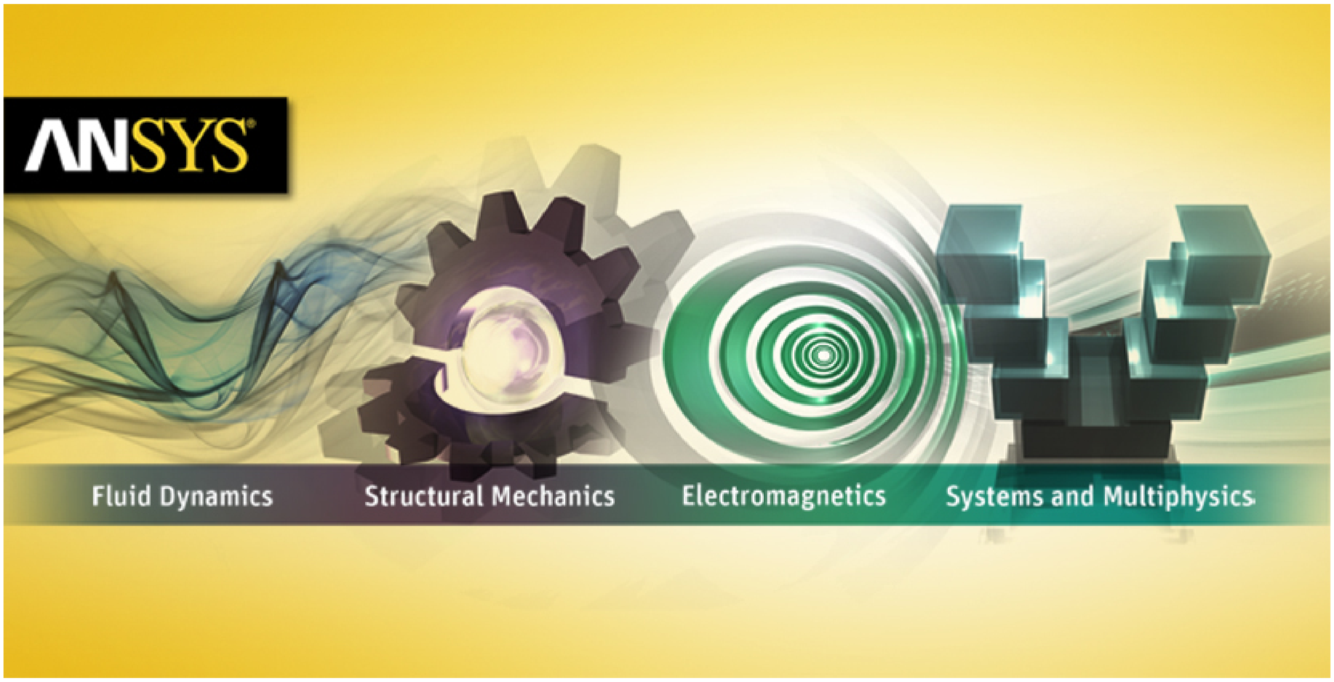


# Running ANSYS FLUENT Under SGE

ANSYS, Inc.
Southpointe
275 Technology Drive
Canonsburg, PA 15317
ansysinfo@ansys.com
http://www.ansys.com
(T) 724-746-3304
(F) 724-514-9494

Release 14.0
November 2011

ANSYS, Inc. is certified to ISO 9001:2008.

Published in the U.S.A.

# Table of Contents

# About This Document

This document provides general information about running ANSYS FLUENT under Sun Grid Engine software (SGE). Examples have also been included, where available.

Information in this document is presented in the following chapters:



This document is made available via the ANSYS, Inc. website for your convenience. Please contact Oracle, Inc. (http://www.oracle.com/) directly for support of their product.

# Chapter 1: Introduction

Oracle (Sun) Grid Engine (SGE) software is a distributed computing resource management tool that you can use with either the serial or the parallel version of ANSYS FLUENT. This document provides general information about running ANSYS FLUENT under SGE, and is made available via the ANSYS, Inc. website for your convenience. Please contact Oracle, Inc. (http://www.oracle.com/) directly for support of their product.

ANSYS FLUENT submits a process to the SGE software, then SGE selects the most suitable machine to process the ANSYS FLUENT simulation. You can configure SGE and select the criteria by which SGE determines the most suitable machine for the ANSYS FLUENT simulation.

Among many other features, running an ANSYS FLUENT simulation using SGE enables you to:

- Save the current status of the job (this is also known as checkpointing when the ANSYS FLUENT `.cas` and `.dat` files are saved)
- Migrate the simulation to another machine
- Restart the simulation on the same or another machine.

For more information, see the following section:


## 1.1. Overview of ANSYS FLUENT and SGE Integration

For more information, see the following sections:


### 1.1.1. Requirements

- Oracle (Sun) Grid Engine software version 6.x, which is available online at http://www.oracle.com/
- FLUENT 6.0–ANSYS FLUENT 14.0

**Important**

Running ANSYS FLUENT under SGE is not supported on Windows.

### 1.1.2. ANSYS FLUENT and SGE Communication

ANSYS FLUENT and SGE communicate with each other through checkpointing and migration commands. To checkpoint, or save, ANSYS FLUENT simulations, SGE uses an executable file called `ckpt_com-`

mand.fluent. To migrate ANSYS FLUENT simulations to another machine, SGE uses another executable file called migr_command.fluent.

### 1.1.3. Checkpointing Directories

ANSYS FLUENT creates a checkpointing subdirectory, identified by the job ID. The checkpointing directory contains files related only to the submitted job.

### 1.1.4. Checkpointing Trigger Files

When an ANSYS FLUENT simulation needs to be checkpointed, SGE calls a command that creates a checkpoint trigger file (check) in the job subdirectory, which causes ANSYS FLUENT to checkpoint and continue running. If the job needs to be migrated, because of a machine crash or for some other reason, a different trigger file (exit) is created which causes ANSYS FLUENT to checkpoint and exit.

### 1.1.5. Default File Location

For ANSYS FLUENT 14.0, SGE-related files are installed by default in *path*/ansys_inc/v140/fluent/fluent14.0.0/addons/sge, where *path* is the ANSYS FLUENT installation directory; for previous versions of ANSYS FLUENT (prior to version 12.0), the location is *path*/Fluent.Inc/addons/sge1.0. The files include the following:

- ckpt_command.fluent
- migr_command.fluent
- sge_request
- kill-fluent
- sample_ckpt_obj
- sample_pe

These files are described in the sections that follow.

# Chapter 2: Configuring SGE for ANSYS FLUENT

SGE must be installed properly if checkpointing is needed or parallel ANSYS FLUENT is being run under SGE. The checkpoint queues must be configured first, and they must be configured by someone with manager or root privileges. The configuration can be performed either through the GUI `qmon` or the text command `qconf`.

For more information, please see the following sections:


## 2.1. General Configuration

Using the SGE graphical interface, the general configuration of SGE and ANSYS FLUENT requires the following:

- The **Shell Start Mode** must be set to either **unix_behavior** or **posix_compliant**.
- Under **Type**, the **checkpointing** option must be marked as true.

When running parallel ANSYS FLUENT simulations, the following options are also important:

- Under **Type**, the **parallel** option must be marked as true.
- The value of **slots** should be set to a value greater than `1`.

## 2.2. Checkpoint Configuration

Checkpointing can be configured using the **min_cpu_interval** field. This field specifies the time interval between checkpoints. The value of **min_cpu_interval** should be a reasonable amount of time. Entering a value that is too low for **min_cpu_interval** results in frequent checkpointing operations, and writing `.cas` and `.dat` files can be computationally expensive.

SGE requires checkpointing objects to perform checkpointing operations. ANSYS FLUENT provides a sample checkpointing object called `sample_ckpt_obj`.

Checkpoint configuration also requires root or manager privileges. While creating new checkpointing objects for ANSYS FLUENT, keep the default values as given in the sample/default object provided by ANSYS FLUENT and change only the following values:

- queue list (`queue_list`)

  The queue list should contain the queues that are able to be used as checkpoint objects.

- checkpointing and migration commands (`ckpt_command` and `migr_command`)

  These values should only be changed when the executable files are not in the default location, in which case the full path should be specified. All the files (that is, `ckpt_command.fluent` and migr_command.fluent) should be located in a directory that is accessible from all machines where the ANSYS FLUENT simulation is running. When running ANSYS FLUENT 14.0, the default location for these files is *path*/ansys_inc/v140/fluent/fluent14.0.0/addons/sge, where *path* is the ANSYS FLUENT installation directory; for previous versions of ANSYS FLUENT (prior to version 12.0), the location is *path*/Fluent.Inc/addons/sge1.0.

- checkpointing directory (ckpt_dir)

  This value dictates where the checkpointing subdirectories are created, and hence users must have the correct permission to this directory. Also, this directory should be visible to all machines where the ANSYS FLUENT simulation is running. The default value is NONE where ANSYS FLUENT uses the current working directory as the checkpointing directory.

- checkpointing modes (when)

  This value dictates when checkpoints are expected to be generated. Valid values of this parameter are composed of the letters s, m, and x, in any order:

  – Including s causes a job to be checkpointed, aborted, and migrated when the corresponding SGE **Exceed** daemon is shut down.

  – Including m results in the generation of checkpoints periodically at the *min_cpu_interval* interval defined by the queue (see qconf).

  – Including x causes a job to be checkpointed, aborted, and migrated when a job is suspended.

**Important**

The m mode must be set to permit interval checkpointing.

## 2.3. Configuring Parallel Environments

For submitting parallel jobs, SGE needs a parallel environment (PE) interface. ANSYS FLUENT provides a sample parallel environment interface called fluent_pe.

Parallel environment configuration requires root or manager privileges. Change only the following values when creating a new parallel environment for ANSYS FLUENT:

- queue list (queue_list)

  This should contain all the queues where qtype has been set to PARALLEL.

- user/xuser lists (user_list and xuser_lists)

  These contain the lists of users who are allowed or denied access to the parallel environment.

- Shut-down procedure invocation command (stop_proc_args)

  This should be changed only if the kill-fluent executable is not in the default directory, in which case the full path to the file should be given and the path should be accessible from every machine.

- slots (slots)

  This should be set to a large numerical value, indicating the maximum of slots that can be occupied by all the parallel environment jobs that are running.

Since ANSYS FLUENT uses `fluent_pe` as the default parallel environment, an administrator must define a parallel environment with this name.

## 2.4. Default Request File

As an alternative to specifying numerous command line options or arguments when you invoke SGE, you can provide a list of SGE options in a default request file. All default request options and arguments for SGE that are common to all users can be placed in this file.

A default request file should be set up when using SGE with ANSYS FLUENT. ANSYS FLUENT provides a sample request file called `sge_request`. To learn more about how to configure and utilize this resource, see the relevant documentation available at www.oracle.com.

Individual users can set their own default arguments and options in a private general request file called `.sge_request`, located in their `$HOME` directory. Private general request files override the options set by the global `sge_request` file.

Any settings found in either the global or private default request file can be overridden by specifying new options in the command line.

# Chapter 3: Running an ANSYS FLUENT Simulation under SGE

Information in this chapter is divided into the following sections:



## 3.1. Submitting an ANSYS FLUENT Job from the Command Line

When submitting an ANSYS FLUENT job from the Linux line command using SGE, you must include additional options for the ANSYS FLUENT command syntax. The command line syntax is as follows:

```
fluent solver_version [FLUENT_options] -sge [-sgeckpt ckpt_object]
[-sgeq queue_name] [-sgepe parallel_env MIN_N-MAX_N]
```

where:

- *solver_version* specifies the dimensionality of the problem and the precision of the ANSYS FLUENT calculation (for example, `3d`, `2ddp`).
- *FLUENT_options* specify the startup option(s) for ANSYS FLUENT, including the options for running ANSYS FLUENT in parallel. For more information, see the ANSYS FLUENT User's Guide.
- `-sge` is a start up option that tells ANSYS FLUENT to run under SGE.
- `-sgeckpt` *ckpt_object* is a startup option that specifies the checkpointing object and overrides the checkpointing option specified in the default request file. If this option is not specified in the command line and the default general request file contains no setting, then the ANSYS FLUENT simulation is unable to use checkpoints.
- `-sgeq` *queue_name* is a startup option that specifies the name of the queue.
- `-sgepe` *parallel_env MIN_N-MAX_N* is a startup option that specifies the parallel environment to be used when ANSYS FLUENT is run in parallel. This should be used only if the `-sge` option is used.

  The specified *parallel_env* must be defined by an administrator. For more information about creating a parallel environment, refer to the `sample_pe` file that is located in the `/addons/sge` directory in your ANSYS FLUENT installation area.

  The values for *MIN_N* and *MAX_N* specify the minimum and maximum number of compute nodes, respectively.

  If ANSYS FLUENT is run in parallel under SGE and the `-sgepe` parameter is not specified, by default it will attempt to utilize a parallel environment called `fluent_pe`. Note that `fluent_pe` must be defined by an administrator if you are to use this default parallel environment. In such a case, *MIN_N* will be set to `1` and *MAX_N* will be set to the maximum number of requested compute nodes specified in the *FLUENT_options* (for example, `-t4`).

The following examples demonstrate some applications of the command line syntax:

- Serial 2D ANSYS FLUENT simulation running under SGE

  ```
  fluent 2d -sge
  ```

- Serial 2D ANSYS FLUENT simulation with checkpoints, in which `fluent_ckpt` is the checkpointing object

```
fluent 2d -sge -sgeckpt fluent_ckpt
```

> **Important**
>
> You can use `qconf -sckptl` to list available checkpoint objects.

- Parallel 2D ANSYS FLUENT simulation running under SGE on 4 CPUs

```
fluent 2d -t4 -sge
```

- Parallel 2D ANSYS FLUENT under SGE on a range of nodes (between 2 and 6), using the parallel environment `diff_pe`

```
fluent 2d -t4 -sge -sgepe diff_pe 2-6
```

> **Important**
>
> In this example, note that the `-t4` option will be ignored and the `2-6` option will take precedence.

- Parallel 2D ANSYS FLUENT under SGE on 4 CPUs, using the parallel environment `diff_pe` and the queue `large`

```
fluent2d -t4 -sge -sgeq large -sgepe diff_pe 4 -4
```

## 3.2. Submitting an ANSYS FLUENT Job Using FLUENT Launcher

FLUENT Launcher has graphical user input options that allow you to submit an ANSYS FLUENT job using SGE. Perform the following steps:

1. Open FLUENT Launcher (*Figure 3.1* (p. 9)) by entering `fluent` without any arguments in the Linux command line.

**Figure 3.1 The Scheduler Tab of FLUENT Launcher (Linux Version)**

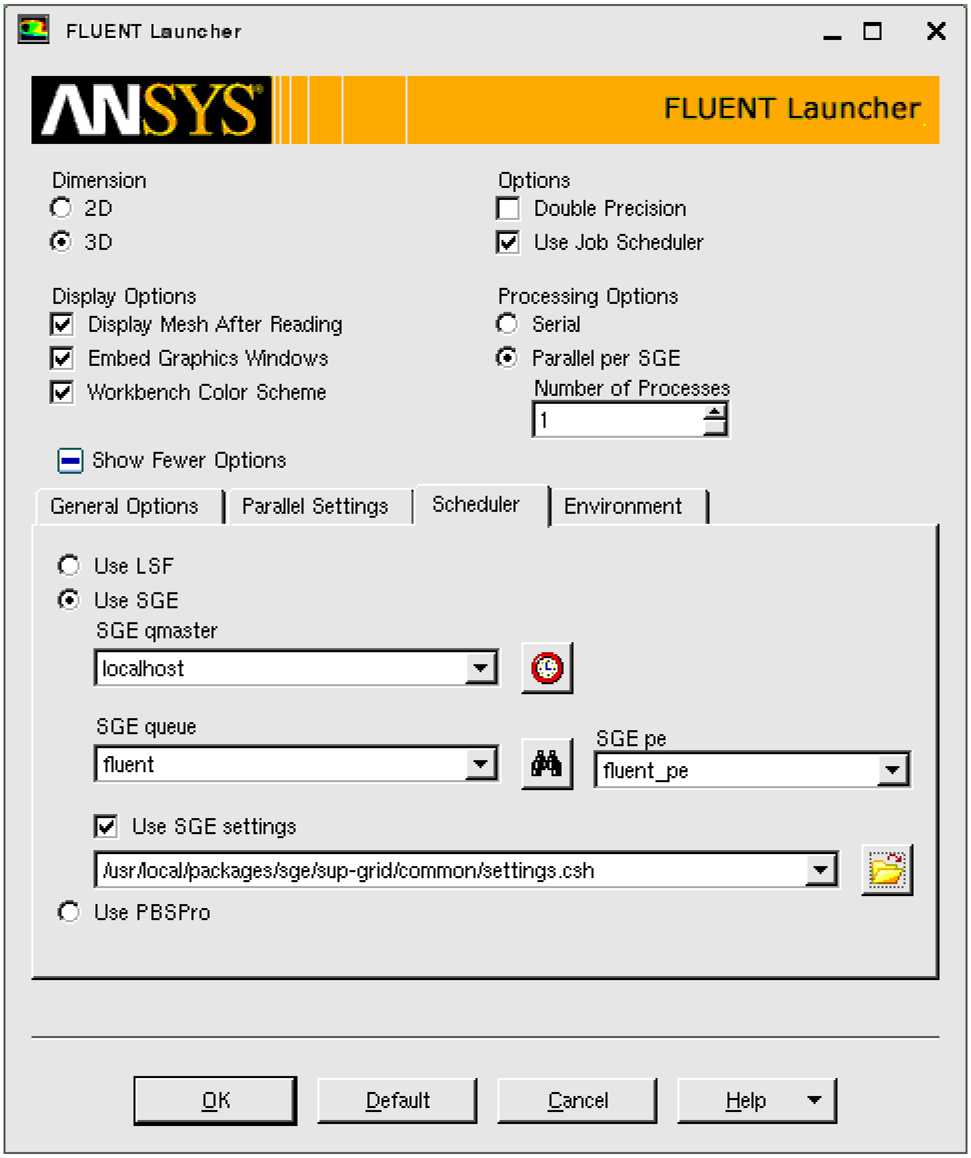


2. Enable **Use Job Scheduler** under **Options**.
3. Click **Show More Options** to expand FLUENT Launcher.
4. Click the **Scheduler** tab.
    a. Select **Use SGE**.
    b. Enter the name of a node for **SGE qmaster**. SGE will allow this node to summon jobs. By default, **local host** is specified for **SGE qmaster**. Note that the  button enables you to check the job status.
    c. You have the option of entering the name of a queue in which you want your ANSYS FLUENT job submitted for **SGE queue**. Note that you can use the  button to contact the **SGE qmaster** for a list of queues.
    d. If you are running a parallel simulation, you must enter the name of the parallel environment in which you want your ANSYS FLUENT job submitted for **SGE pe**. The parallel environment must

be defined by an administrator. For more information about creating a parallel environment, refer to the `sample_pe` file that is located in the `/addons/sge` directory in your ANSYS FLUENT installation area.

   e. You can specify an SGE configuration file by enabling the **Use SGE settings** option. Then enter the name and location of the file in the text box or browse to the file.

5. Set up the other aspects of your ANSYS FLUENT simulation using the FLUENT Launcher GUI items. For more information, see the ANSYS FLUENT User's Guide.

# Chapter 4: Running ANSYS FLUENT Utility Scripts under SGE

You can run the ANSYS FLUENT utility scripts (for example, `fe2ram`, `fl42seg`, `ic3m2ram`, `partition`, `tconv`, `tmerge`, `viewfac`, `tpoly`) with the SGE load management system.

The command line syntax to launch an ANSYS FLUENT utility script under SGE is as follows:

```
utility utility_name [utility_options] -sge [-sgeq queue_name] [-sgepe parallel_env MIN_N-MAX_N] utility_inputs
```

where

- *utility_name* is the name of the utility to be launched (for example, `fe2ram`, `fl42seg`, `ic3m2ram`, `partition`, `tconv`, `tmerge`, `viewfac`, `tpoly`).
- *utility_options* are the options that are part of the syntax of the utility being launched. For more information about the options for the various utilities, see the ANSYS FLUENT User's Guide.
- `-sge` is a start up option that tells `utility_name` to run under SGE.
- `-sgeq` *queue_name* is a startup option that specifies the name of the queue.
- `-sgepe` *parallel_env MIN_N-MAX_N* is a startup option that specifies the parallel environment to be used when *utility_name* is run in parallel. This should be used only if the `-sge` option is used.

  The specified *parallel_env* must be defined by an administrator. For more information about creating a parallel environment, refer to the `sample_pe` file that is located in the `/addons/sge` directory in your ANSYS FLUENT installation area.

  The values for *MIN_N* and *MAX_N* specify the minimum and maximum number of compute nodes, respectively.

  If *utility_name* is run in parallel under SGE and the `-sgepe` parameter is not specified, by default the utility will attempt to utilize a parallel environment called `utility_pe`. Note that `utility_pe` must be defined by an administrator if you are to use this default parallel environment. In such a case, *MIN_N* will be set to `1` and *MAX_N* will be set to the maximum number of requested compute nodes specified in the *utility_options* (for example, `-t4`).

  For example, to run the `viewfac` utility on `file_name` in serial, the line command would be the following:

  ```
  utility viewfac -sge file_name
  ```

  In parallel with 4 CPUs, without the use of `-sgepe` (and therefore using the parallel environment `utility_pe`), the command would be:

  ```
  utility viewfac -t4 -sge file_name
  ```

  In parallel on a range of nodes (between 1 and 2) with the use of `-sgepe` and the parallel environment `diff_pe`, the command would be:

  ```
  utility viewfac -t4 -sge -sgepe diff_pe1 -2 file_name
  ```

- *utility_inputs* are the fields that are part of the syntax of the utility being launched. For more information about the necessary fields for the various utilities, see the ANSYS FLUENT User's Guide.

---

**Important**

Note that neither checkpointing, restarting, nor migrating are available when using the utility scripts under SGE.

For more information, please see the following section:


## 4.1. Special Considerations for Running the viewfac Utility under SGE

Given a cluster file as input, the `viewfac` utility is provided for the computation of viewfactors (required in ANSYS FLUENT when using the S2S radiation model). Using SGE, you can create a cluster file in ANSYS FLUENT and choose to compute and write viewfactors in the same ANSYS FLUENT session, or you can manually launch the `viewfac` utility from outside ANSYS FLUENT and compute the viewfactors. When viewfactors are computed from ANSYS FLUENT, the following must be considered:

- If ANSYS FLUENT is launched without the `-sge` option, then the `viewfac` utility should also be launched from ANSYS FLUENT without the `-sge` option.
- If ANSYS FLUENT is launched under SGE (that is, by using the `-sge` option), then you should be able to launch the `viewfac` utility either under SGE or without SGE. By default, the utility is launched under SGE, with the same SGE related parameters as those used to launch ANSYS FLUENT under SGE.

---

**Important**

To launch a parallel version of the `viewfac` utility, the parallel environment `utility_pe` must be set up and configured in a manner that is suitable for your jobs. You may need to contact your SGE administrator for assistance.

If you do not wish to launch the utility under SGE, then you may use the Scheme functions described as follows.

The Scheme function `(sge?)` can be used to inquire whether the `viewfac` utility is enabled or disabled under SGE. The function will return either `#t` (true) or `#f` (false).

To switch the value of the `(sge?)` function, use the `(switch-sge)` Scheme function. A message will state the new value (for example, `"SGE enabling set to #t"`). Again, by default, the `viewfac` utility will use the same SGE related options as ANSYS FLUENT. The SGE related options can also be controlled through the `(set-sge-optionsstring_arg)`Scheme function. The `string_arg` argument replaces the actual `sge-options` that you originally set (for example, `(set-sge-options "-q dev-test")`). Note that you should take care to set the SGE options correctly if you use the `(set-sge-options string_arg)` function, because this string is simply passed to SGE in its command and is interpreted by SGE (and not by the utility script).

# Index