




Scientific Computing & Modelling

# Installation Manual

## ADF Program System Release 2008.01

Scientific Computing & Modelling NV  
Vrije Universiteit, Theoretical Chemistry  
De Boelelaan 1083; 1081 HV Amsterdam; The Netherlands  
E-mail: [support@scm.com](mailto:support@scm.com)

Copyright © 1993-2008: SCM / Vrije Universiteit, Theoretical Chemistry, Amsterdam, The Netherlands  
All rights reserved



# Table of Contents

Installation Manual .....	1
Table of Contents .....	2
1 Introduction .....	3
1.1 Requirements .....	3
2 Installation.....	5
2.1 Decide which version to install .....	5
2.2 Download the software .....	5
2.3 Install external libraries (optional) .....	6
2.4 Set up your environment .....	7
2.5 Source distribution only: unpack, configure and compile the source .....	9
2.6 Generate license information .....	10
2.6.1 Floating License.....	12
2.7 Setup the nodeinfo file (PVM only) .....	13
2.8 Set up scratch space.....	14
2.9 Run some examples and test performance .....	15
2.10 Clean up.....	16
3 Directory structure of the ADF package .....	17
3.1 The bin directory.....	17
3.2 The atomicdata directory .....	17
3.3 The examples directory.....	17
3.4 The source code (adf, band, libtc, Install, ...) .....	17
3.5 The Doc directory .....	18
3.6 The scripting directory.....	18
4 Additional Information .....	19
4.1 Performance-related environment variables.....	19
4.2 MPI .....	20
4.3 More on HP-MPI .....	21
4.4 PVM.....	22
4.5 Known HP-MPI/Linux issue: libibverbs error message .....	22

# 1 Introduction

The Installation Manual describes the installation of the ADF program package on the platforms on which it is supported. See also the [ADF installation manual for Windows](#) and the [ADF installation manual for Mac OS X](#). For optimal performance, some system specific tuning may be needed. We therefore strongly recommend to also read Chapter 4 of this document. This Installation Manual is available from the Documentation part of the SCM web site:

<http://www.scm.com/Doc/Welcome.html>

The main programs contained within the ADF package are ADF and BAND. There are also several utility and property programs. These are used primarily for pre- and post-processing data of the main programs. You will always install the complete ADF package at once. It includes ADF, BAND, and all available utility programs, as well as all of the ADF-GUI modules.

The ADF package and support scripts are written for use within a UNIX environment (including Linux). If you want to install it and use it effectively you will need some basic UNIX knowledge. We do **not** provide this general UNIX introduction. If you are new to UNIX we strongly suggest you first make yourself comfortable with this operating system. Very nice tutorials and introductions are available for free on the web, simply search for 'UNIX introduction' or 'UNIX tutorial'.

Typically installation of the ADF package is simple and straightforward. If you have problems installing it, contact us for assistance at [support@scm.com](mailto:support@scm.com).

## 1.1 Requirements

The requirements for the ADF package vary from platform to platform. A list of supported platforms, with information on the operating system (versions), parallel environment (MPI, PVM, ...), and compilers used is available on the Download section of our web site:

<http://www.scm.com/Downloads/Welcome.html>

### Memory

The amount of memory you need depends greatly on the kind of calculation you perform. For most calculations 256 MB per CPU will be sufficient, but if you have lots of memory available most program will run significantly faster. It also reduces the amount of disk I/O, which may speed up your calculation depending on the I/O sub-system and the operating system.

Memory requirement grows with the system size. For example, for a molecule containing 100 atoms with a DZ basis set it may be sufficient to have 256 MB per CPU but for a molecule with 1000 atoms up to 8 **gigabytes** per CPU may be required.

### Disk

For the installation of the package you need from less than 256 MB (without sources) to 1 GB (with sources, and compiled objects). On Windows you will need about 550MB of disk space for ADF installation. The (scratch) disk space requirements for running ADF greatly depend on what type of calculation you perform. For most ADF calculations 1 GB of disk space will be enough. For BAND calculations you may need from a few to up to a few tens of gigabytes (for large calculations) of free disk space.

### Operating System

The package runs on most recent UNIX operating systems, including Linux.

A Windows version is also available. It has been tested under Windows XP and Vista. See the [ADF installation manual for Windows](#) for more details.

For Macintosh OS X you will need at least 10.3 (Panther), and you need to install the X11 environment as well (it is available as an optional install on the Panther installation disks).

If your UNIX platform is not listed, please contact us for more information.

### **Network**

First of all, on the most popular systems (Linux, MacOS and Windows) a network card must be present in the computer as its hardware MAC address is used as the computer's ID for licensing purpose. Besides that, when using HP-MPI, the computer must be connected to a network and have a valid IP address and the address must not change while an MPI job is running. This requirement might be lifted in a future version of HP-MPI.

As far as performance concerned, a switched (Gigabit or Fast) Ethernet network is typically sufficient for good parallel performance on up to eight nodes. If you are using more nodes you may need a faster communication method, such as Myrinet or Infiniband, to get the best performance.

All commercial parallel machines (SGI, IBM, HP, ...) have very good communication facilities, so the package should run fine on those platforms as far as the communication system is concerned.

### **Compiler**

If you have a license for the source code, you can compile and link the source yourself, with or without your own modifications.

The source consists mainly of Fortran95 code, with some small parts written in C. You need to use the same Fortran and C compilers as we are using on the same platform, since some parts of the code are available only as object modules. It is very unlikely that other compilers will work with these object modules. We cannot support compilers other than what we are using ourselves.

To check what compiler to use, check the detailed machine information on the Download section of our web site.

## 2 Installation

To install the ADF package you have to go through the following steps:

- 1. Decide which version to install (serial or MPI)
- 2. Download and unpack the software
- 3. Install external libraries (optional)
- 4. Set up your environment
- 5. Source distribution only: compile the source
- 6. Generate license information
- 7. Setup a nodeinfo file (PVM only)
- 8. Setup scratch space
- 9. Run some tests
- 9. Clean up

Below we discuss each step separately.

### 2.1 Decide which version to install

If you are installing the package on a machine that has only one CPU you should install the **serial version**. For some platforms we have no serial version available. In such cases you can just install any parallel version.

If you are installing the package on a machine with more than one CPU or on a multi-core machine, or if you are installing ADF on a cluster, you should install a **parallel version**. Please note that the Hyperthreading technique available in some processors does not improve ADF performance. Thus, even though a hyperthreaded processor is seen by the system as two logical processors, it should be considered as one.

For most systems, only an MPI version is provided but for some systems there is also a serial version. If there is no serial version available for your system then you can run an MPI version in serial mode.

Note about PVM. Even though PVM is still supported by the ADF source code, SCM does not distribute any PVM versions of ADF. Thus, only those users who compile ADF themselves have access to the PVM version.

### 2.2 Download the software

If you have not already done so, you need to download the software from the Download section of the SCM web site. You will be asked for a username and password that you should have received from SCM. Contact [admin@scm.com](mailto:admin@scm.com) if you forgot your username and/or password.

You need to download the binary tarball for your platform in all cases, also if you want to compile ADF from sources. For example, for 32-bit Linux with HP-MPI the file is called `adf2008.01.pc_linux.hpmpi.bin.tgz`. If you have a source code license and want to compile the source code yourself you also need to download `adf2008.01.src.tgz`.

- Download the `adf2008.01.bin.tgz` and, if applicable, `adf2008.01.src.tgz` for your platform (from [www.scm.com](http://www.scm.com)).

**Important:** skip the step below if you have a source license and are going compile ADF yourself and move on to the next section.

If you want to install the precompiled executables only then gunzip and untar the downloaded files:

```
gunzip adf2008.01.bin.tgz
tar -xf adf2008.01.bin.tar
```

A new directory adf2008.01 will be created containing your new adf installation.

## 2.3 Install external libraries (optional)

If you have downloaded an HP-MPI or OpenMPI version of ADF you do not need to install any additional MPI libraries because these are provided by SCM.

If, however, for whatever reason you want to build ADF with a different MPI library then below you will find some information and web links we have gathered about major MPI implementations supported by ADF. Additional information about different MPI libraries, caveats and troubleshooting may be found in the section 4 of this document.

### HP-MPI

Hewlett-Packard has MPI implementations for the most popular operating systems: x86- and ia64-based Linux, Windows, as well as for HP-UX. Full documentation about HP-MPI is available at the [HP website](#). HP-MPI is distributed with (HP-MPI versions of) ADF and a complete HP-MPI directory tree is found in \$ADFBIN/hpmi. Using HP-MPI with ADF does not require any additional licenses. In addition to TCP/IP and shared memory supported by every MPI implementation, HP-MPI also supports (without recompilation) all well-known interconnects such as Infiniband and Myrinet. The best available interconnect is chosen automatically at run time and this can also be modified using mpirun command-line switches in the \$ADFBIN/start script.

### OpenMPI

OpenMPI is an open-source implementation of the MPI standard. You can get both the software and documentation from the OpenMPI website:

<http://www.openmpi.org>

For some platforms, ADF comes bundled with OpenMPI so you do not have to install it separately. However, you **do** need to install OpenMPI if you are going to recompile ADF and if you want it to use a fast interconnect such as Infiniband or if you need to enable OpenMPI integration with the batch system on your cluster.

### Intel MPI

Intel provides its own MPI implementation that is similar to OpenMPI in many ways. ADF supports Intel MPI at the configure script level on ia32-, i64- and ia64-based Linux platforms.

### MPICH1

MPICH1 is a portable, though outdated, implementation of the MPI v1 standard. You can get both the software and documentation from the MPICH1 web site:

<http://www-unix.mcs.anl.gov/mpi/mpich1/>

### Vendor-specific MPI implementations

Many hardware manufacturers, such as SGI, IBM, Sun and Cray, have their own MPI implementations that work only on their hardware and/or operating system. In this case, SCM provides an ADF version built with the appropriate MPI library but the run-time environment (RTE) is not distributed since you will most likely

have the RTE already pre-installed on your machine. Please contact the machine vendor's support staff if this is not the case.

## PVM

**NOTE:** SCM strongly discourages the use of PVM. It is not actively developed, it does not use modern fast interconnects and it is significantly slower than any modern MPI implementation. So it should only be used as the last resort when no other parallel option works.

If you are not familiar with PVM and need to install it yourself, we strongly suggest you have a look at the PVM documentation. PVM (including FAQ, User's Guide and the Tutorial) can be downloaded from:

<http://www.netlib.org/pvm3/index.html>

To use PVM you need to:

- install PVM (see PVM documentation)
- set two environment variables: PVM\_ROOT and PVM\_ARCH
- add \$PVM\_ROOT/bin/\$PVM\_ARCH and \$PVM\_ROOT/lib to your PATH
- make sure pvm can start the PVM daemon pvmd on all nodes of your system

See section 4 for more information about PVM troubleshooting

**Note:** you should use PVM\_ARCH=**MACOSX** on the Macintosh OS X platform. If you choose something different you will certainly encounter difficulties.

## 2.4 Set up your environment

You will need to define a couple of environment variables, preferably in your login script.

If you are using a of batch system, such as PBS, SGE or LSF, then make sure that the variables are set not only for interactive work but also for batch use. In a distributed parallel environment, like a Linux cluster, make sure that these environment variables are set on all nodes.

The most important environment variables are automatically passed on from the master process to the child processes in case of a parallel run. Thus the values set on the node where you start the program will override any value you may have set on the nodes where the child processes are running. Typically, you should still set the environment variables on all nodes since you won't know in advance on which node the master will run.

Hint: set the SCM\_DEBUG environment variable to "yes" before running ADF to see values of relevant environment variables for each process. This may be especially useful if you suspect that some variables are not propagated through mpirun command to the master node.

The following environment variables must always be set for all ADF versions.

Name	Possible value	Description
ADFHOME	\$HOME/adf2008.01	full path of the ADF installation directory (see section 2.3)
ADFBIN	\$ADFHOME/bin	full path ADF directory for scripts and binaries
ADFRESOURCES	\$ADFHOME/atomicdata	full path of the directory with the ADF database (basis sets and so on)
SCMLICENSE	\$ADFHOME/license.txt	full path to the license file

The following environment variables may be required for any parallel (MPI or PVM) version.

Name	Possible value	Description
NSCM	16	default number of parallel processes; do <b>not</b> put it in the shells's rc file but rather set this variable per job
SCM_TMPDIR	/scratch/\$USER	full path of the directory where all processes will create their temporary files. See also sections 2.8 and 4.

The following environment variables are relevant for MPI versions only.

Name	Possible value	Description
SCM_MACHINEFILE	\$HOME/.machine	full path-name of the file containing a list nodes to use for a job; <b>Important:</b> this variable should be set <b>per job</b> and it should <b>only</b> be set if multiple computers are used without a batch system. The file pointed to by the variable must already exist and it must contain a list of nodes on which a particular job is about to be executed, with their processor count. See section 4.2 for more information.

The following environment variables are relevant for a PVM version only.

Name	Possible value	Description
PVM_ROOT	\$HOME/pvm3	full path of the main PVM installation directory
PVM_ARCH	SGI6	architecture code for your machine (NOTE: use SGI6 on SGI IRIX)
SCMSPAWNSCRIPT	\$ADFBIN/adfs	full path of the script used by the PVM daemon to start slave nodes

The following environment variables are relevant for the GUI modules

Name	Default Value	Description
SCM_ERROR_MAIL	no error e-mail sent	e-mail address for error reports
SCM_GUIRC	\$HOME/.scm_guiirc	location of the preferences file
SCM_TPLDIR	none (no extra templates loaded)	location of the templates directory
SCM_STRUCTURES	none (no extra structures loaded)	location of the structures directory
SCM_RESULTDIR	none (current directory used)	location of the results directory
DISPLAY	required (for all X11 programs) except on Windows	X-window display to use
SCM_MOPAC	none (default script \$ADFBIN/mopac.scm will be used)	command to start MOPAC
SCM_QUEUES	none (ADFjobs will search the remote \$HOME/.scmgui)	path to the dynamic queues directory

The following environment variables are relevant for source distributions only, and only at the configure time.

Name	Possible value	Description
------	----------------	-------------

MPIDIR	/opt/hpmpi	full path of the directory where MPI library is installed. This variable is required at configure time when a non-standard MPI library is used. <b>Important:</b> Do <b>not</b> set MPIDIR when building an HP-MPI or OpenMPI version of ADF with MPI libraries delivered with ADF.
MATHDIR	/opt/intel/mkl/10.0.011	(optional) installation directory of the Intel MKL library

The following environment variables may be set to modify other aspects of ADF execution. All of them are optional and some are used for debugging only.

Name	Possible value	Description
SCM_IOBUFFERSIZE	64	Size of memory buffer for scratch files, in megabytes. See also section 4.
SCM_VECTORLENGTH	128	Block length used in the numerical integration routines. See also section 4
SCM_RCV_TIMEOUT	900	(PVM only) time (seconds) to wait for messages
SCM_MAX_RCV_TIMEOUTS	2	(PVM only) maximum number of timeouts
SCM_ALARMTIME	1800	Time (seconds) to wait for messages using ALARM signal, 0 disables timeouts. Currently the default setting is zero, because this feature may not work on some platforms.
SCMB COP	1	Broadcast algorithm (not to be changed for MPI versions)
SCMGAOP	5	Gather algorithm (not to be changed for MPI versions)
SCMCBOP	1	Combine algorithm (not to be changed for MPI versions)
SCM_DEBUG	yes	Set to debug SCM_TMPDIR usage and environment variables (also see above)
SCM_NOMEMCHECK	yes	Disable memory checks
SCM_NODOMAINCHECK	yes	Disable domain check for license
SCM_TRACETIMER	yes	For debugging only: produce additional output on ADF internal timer entry/exit
SCM_DEBUG_ALL	yes	For debugging only: equivalent to specifying DEBUG \$ALL in the input

## 2.5 Source distribution only: unpack, configure and compile the source

The downloaded sources should be unpacked first (i.e. before unpacking binaries), for example, as follows:

```
gzip -cd adf2008.01.src.tgz | tar xf -
```

The result will be a new adf2008.01 directory containing the sources.

Next, unpack the binary distribution. This installs a couple of programs in the \$ADFBIN directory that are not included in the source distribution.

```
gzip -cd adf2008.01.bin.tgz | tar xf -
```

After unpacking everything and setting up your environment properly (see previous section), you should run the configure script. This script is located in the \$ADFHOME/Install directory, and it must be executed from the \$ADFHOME directory. The script replaces some files in the bin directory with versions specifically targeted for your system. For example, if you want to install an OpenMPI version of ADF on a platform for which only an HP-MPI version is available, then *configure* will make an \$ADFBIN/start script suitable for OpenMPI and overwrite the old one which was made for HP-MPI. Further, *configure* creates the Makeflags and settings files that you will need to compile ADF.

Example:

```
cd $ADFHOME
Install/configure -p hpmi
```

The configure script accepts the following options:

- **-h** with this option configure will print the system name, processor type and, for Linux, the GLIBC version. It will also print the default machine name used internally by ADF and a list of known parallel implementations for this combination, as well as a list of possible optional parameters, if any (see below).
- **-p par** specifies the parallel implementation to use. The list of possible **par** values varies per platform and can be obtained by running configure with the **-h** option. The default is "serial".
- **-n ncpus** specifies the maximum number of processes on which you can later run ADF and other parallel programs from the package. The default value of 128 is sufficient in most cases.
- **-o option** specifies an optional parameter if available. For example, on IBM AiX you can specify **-o 64** to configure a version with 64-bit addressing.

Next, you can compile and link the ADF sources by executing the yam (Yet Another Make) script created in the \$ADFBIN directory by configure.

```
cd $ADFHOME
bin/yam
```

After a while, depending on the speed of your computer, ADF should be ready to use, just as if you had installed the precompiled executables.

## 2.6 Generate license information

Most of the programs that form the ADF package require a license file to run. This file contains information about your machine, the version(s) of the package for which you have a license, and the end date of your license.

If you are a trial user, then you should have already received a valid 'demo' license file by e-mail. If you install ADF from a CD, then you should be able to find a demo license file on the CD. In most other cases you will receive a host-locked license file from SCM, but SCM first needs to get some information from you about your machines. You can generate this information by running a program from the package on all machines where ADF is installed as follows:

```
$ADFBIN/dirac check
```

This will show some basic information about your license such as the date until which it is valid and list of packages it is valid for with their respective version numbers. It will also show you up to how many processes can be in an ADF parallel job. An explanatory error message will be printed if there is no valid license found. After that, information about the computer's ID is displayed. This information is necessary for SCM to generate a license file for this particular computer. To obtain only the information needed for generating license file you can use the following command:

```
$ADFBIN/dirac info
```

Thus, in order to obtain a license file it is sufficient to send to [license@scm.com](mailto:license@scm.com) output of the `$ADFBIN/dirac info` command but in case of problems with the license one should send output of the `$ADFBIN/dirac check` command together with the support request.

After receiving the machine ID(s) SCM will prepare a license file matching your license conditions and e-mail it to you with instructions on how to install it.

After receiving a license file you will need to save it to an appropriate location on your computer (usually in the directory where you installed ADF) and set the environment variable `$SCMLICENSE` pointing to the file. For example, if you saved the license file named `license.txt` in `/opt/ADF` then `SCMLICENSE` must be set to `/opt/ADF/license.txt`. Add the `SCMLICENSE` definition to your login script.

Ensure that permissions on the license file allow read access for everyone who is going to run ADF.

After installing the license you should make sure everything is fine by running:

```
$ADFBIN/dirac check
```

This will produce the output similar to the following:

```
Checked: /home/yakovlev/license
```

```
License termination date (mm/dd/yyyy): 7/10/2009
```

According to it, you are allowed to run:

```
ADF version 2008.99
BAND version 2008.99
ADFGUI version 2008.99
BANDGUI version 2008.99
GUI version 2008.99
Utils version 2008.99
NBO version 2008.99
CRS version 2008.99
```

Number of procs you are allowed to use for:

```
ADF : 128 procs
BAND : 128 procs
```

```
=====
release: 2008.01
:chem.vu.nl:
:Linuxmycomputer00:15:17:76:EC:5C:
ncores : 4
```

---

LICENCE INFO READY

Please refer to [ADF for Windows release notes](#) for details on how to obtain the license information under Windows.

### **Corrupted License File**

You may find that, after having installed the license file, the program still does not run and prints a message "LICENSE CORRUPT". There are a few possible causes. To explain how this error may come about, and how you overcome it, a few words on license files.

Each license file consists of pairs of lines. The first of each pair is text that states in a human-readable format a couple of typical aspects: A 'feature' that you are allowed to use (for instance 'ADF'), the expiration date, a (maximum) release (version) number of the software and so on. The second line contains the same information in encrypted format: a long string of characters that appear to make little sense. The program reads the license file and checks, with its internal encrypting formulas, that the two lines match. If not, it stops and prints a "LICENSE CORRUPT" message.

So, there are two common reasons why this may happen:

1. You are using a license file for a version of the software other than your executables correspond to. Newer (major) releases may use a different encryption, so that the match in old license files is not recognized anymore. In particular, the "LICENSE CORRUPT" error will arise if you run ADF1999 (or later) with a license file that was generated for ADF2.3 (or earlier). Verify that your license file and executable belong to the same major release.
2. The license file as it has been created has been modified in some way. Sometimes, people inspect it and 'clean it up' a little bit, for instance by removing 'redundant spaces', or by making some other 'improvements'. Any such modification will break the encryption match and lead to the "LICENSE CORRUPT" error. Sometimes the reason lies in the mailing system: if the file contains long lines the mailer may break them into shorter lines. To verify (and correct) this open the license file in a text editor and see if it really consists of pairs of lines as described above. If not, re-unify the broken lines and try again.

You can use the **fixlic** utility to try to fix this automatically. Please be aware that the **fixlic** utility will try to fix the file pointed to by the \$SCMLICENSE environment variable and replace it with the fixed copy. Thus, you need to make a backup of your license file first and you need to have write permissions for it.

```
cp $SCMLICENSE $SCMLICENSE.backup
$ADFBIN/fixlic
```

## **2.6.1 Floating License**

If you have a floating license, you will need to follow the instructions below to make it work. If you do not have a floating license, you may skip this section.

### **Create a floating license directory**

Make a new directory to keep track of the running ADF processes. This directory must be located on a file system that is shared between all nodes on which you want to run ADF with this floating license.

Thus, if \$ADFHOM is shared and you want (for example) to use the name FloatADF to count the processes:

```
mkdir $ADFHOME/FloatADF
chmod 1777 $ADFHOME/FloatADF
```

In this example, we have given all users read, write and execute permissions to this directory. If you wish to restrict this for security reasons, you may do so as long as all ADF users will have read, write and search permission for this directory.

### E-mail us the license information

Send the result of the following commands (using again the name FloatADF as an example) to [license@scm.com](mailto:license@scm.com):

```
ls -lid $ADFHOME/FloatADF
```

Also, execute the command below on each machine (or cluster node) where you want to run ADF (the ADF package must be installed before running this command) and include the output in your e-mail.

```
$ADFBIN/dirac info
```

### Important:

- You need to mail the **exact output** of the above command, thus with a **full** path to the FloatADF directory.
- You should **not** delete, move or rename this directory once you have sent this information to SCM. Doing so will break your floating license.

### Install license file

SCM will send you a license file by e-mail. Its installation is identical to that for non-floating licenses. Copy it to your system (most conveniently also in a shared directory) and make sure all users have read permission to it:

```
cp license.txt $ADFHOME/license.txt
chmod ugo+r $ADFHOME/license.txt
```

Next make sure ADF will use this license file by setting the SCMLICENSE environment variable (Bourne shell syntax is shown):

```
SCMLICENSE=$ADFHOME/license.txt
export SCMLICENSE
```

## 2.7 Setup the nodeinfo file (PVM only)

If you have installed a PVM version of the ADF package you may need to set up the nodeinfo file.

The nodeinfo file tells ADF how many tasks are allowed to be run on each node in the PVM virtual parallel machine.

If you have a cluster of single-CPU workstations you most probably want to generate at most one task per node. Since this is the default, you do not need to do anything special.

If one or more of your nodes have more than one CPU, you probably want to run more than one task on those nodes. Normally you should generate one task per CPU. The nodeinfo file communicates that information to ADF.

The format of the nodeinfo file is:

- one line per node,
- for each line: node name and a number, separated by a space.

Here node name is the name of your node, as shown using the `pvm conf` command. Number is the *maximum* number of tasks you wish to use on that node. Nodes which are not specified will default to one task to be generated at most.

Thus, if you have two SGI workstations (`silly1` and `silly2`) with 4 CPU's each, the contents of your `nodeinfo` file should be (two lines):

```
silly1 4
silly2 4
```

You should put this `nodeinfo` file in `$ADFBIN` and make it readable for all ADF users.

For some special situations it makes sense to have a unique `nodeinfo` file for each calculation (for example when using a batch system with a large machine). For this purpose the programs within the ADF package first look for a local `nodeinfo` file (in the directory in which the program is started), and next it will try to use the `$ADFBIN/nodeinfo` file. If neither of them exists, ADF will assume that it may generate only one task per node.

## 2.8 Set up scratch space

Most programs from the ADF package use disk for temporary data. This data often takes a significant amount of space and is used frequently. To avoid run-time errors and performance loss you may want to make sure that the file system used for temporary files is both big and fast. Use the `SCM_TMPDIR` environment variable to tell the programs where to create their scratch directories. Please note that `SCM_TMPDIR` should always be set. If it is not set then each process will create its own directory in the current working directory where it was started.

You can set the `SCM_TMPDIR` environment variable, for example, as follows:

```
SCM_TMPDIR=/local/scratch/$USER
export SCM_TMPDIR
```

In the next few paragraphs we will explain exactly what this variable does.

### Child processes

Every parallel job consists of one master and one or more slave tasks. Master and slaves behave a bit differently with respect to their scratch directories. Slave processes will always create a directory for their scratch files in `$SCM_TMPDIR`. For efficiency reasons, that directory should reside on a local disk unless you are using very, very fast shared file system for large files. You need write access to that directory, and the file system should have enough free space. Please note that the `SCM_TMPDIR` environment variable will be passed from the master to slaves. Besides, slave tasks also `chdir` to the directory to avoid any risk that shared files are updated by more than one process at the same time.

After the job is finished, slave processes will delete their scratch directories. This can be disabled by setting the `SCM_DEBUG` environment variable to any text, for example, to "yes". In this case the scratch directory and all its contents will be left intact.

Each slave writes its text output to a file called `KidOutput` located in its scratch directory. In case of an error this file will likely contain some sensible error message. To avoid losing the file in such a case, ADF will copy it to the directory where the job was started as `KidOutput__#`, where `#` is the process' rank.

### Master process or serial runs

The master process (which is the only process in a serial run) will also create its temporary files in its own sub-directory of \$SCM\_TMPDIR. There are some exceptions: Some files, such as logfile and TAPE13, will be created in the directory where ADF was started because they are not performance-critical but are convenient to have in the current directory for different reasons. For example, logfile is nice to have in the current directory in order to follow the calculation progress and the TAPE13 is an emergency restart file that can be used if ADF crashes or is killed.

At the end of a calculation, the master will copy all result files from its scratch directory to the directory where it was started.

## 2.9 Run some examples and test performance

Verify that you can now run examples provided with ADF and that they give correct results. We recommend that you consult the Examples document for notes on such comparisons: non-negligible differences do not necessarily indicate an error.

**Note:** the sample runs are complete Bourne shell scripts that should be executed as such, they are **not** input files to be fed into any program.

### Test parallel performance

It is very important to make sure that computer resources are utilized with the maximum efficiency. Therefore, you should check that each ADF job uses all processors/cores allocated to it and that the network is not overloaded with the disk I/O traffic.

**Processor usage:** Typically, when you submit a parallel job to the batch system you have a possibility to specify how many processors per node (ppn) to use. If the batch system you are using allows this, then make sure that you request ppn equal to the number of physical cores on each node. The latter is usually equal to the number of logical CPUs reported by the system, except when hyperthreading is turned on. With the hyperthreading on, you should set ppn to half of the number of logical processors.

It is also possible that your batch system does not allow you to specify ppn but instead it always assumes that there only one processor per node. In this case, you will need to edit the \$ADFBIN/start file and add some commands for processing the hostfile.

So, now it is time to see that everything is working as intended. To this end, create a reasonably small ADF job and submit it, possibly to run on more than one node. After the job has finished, open the job's .out file and find the table that looks like the following:

Kid Node	Speed	Task ID	Status
0 compute-0-2	-1	0	
1 compute-0-2	-1	1	
2 compute-0-2	-1	2	
3 compute-0-2	-1	3	
4 compute-0-3	-1	4	
5 compute-0-3	-1	5	
6 compute-0-3	-1	6	
7 compute-0-3	-1	7	

Check the names in the "Node" column and verify that the number of tasks per node is correct. If this looks as expected then move to the end of the file to the table called "Timing Statistics". Before the table, there is a line of text beginning with "Total Used" that might look as follows:

Total Used : CPU= 8064.87 System= 1144.31 Elapsed= 9212.99

This shows how much time (in seconds) was spent in the ADF code (CPU) and in the kernel (System), and how long did it take for the calculation to complete (Elapsed). Ideally, the system time should be small compared the CPU time and the latter should be close to the Elapsed time. The system time will not always be a small number but a sum of the System and CPU time should always give a number very close to the Elapsed time. If this is not the case then it means that ADF has to spend a lot of time waiting for something. This can be, for example, disk I/O or network communication. Note: with HP-MPI the network communication time contributes to the System time while with other MPI implementations it may contribute to the difference between Elapsed and the CPU+System time.

If you notice that the system time portion is enormously large or that there is a large difference between CPU+System and the Elapsed time, then repeat the job with the "PRINT TimingDetail" keyword in the input contact SCM support.

## 2.10 Clean up

At this point you may clean up the \$ADFHOMe directory a little.

You may remove the \*.tar files if you have not already done so, and if you are short on disk space you may remove the examples directories. We do advise, however, to keep them around since they contain useful examples that may be used to get acquainted with ADF.

## 3 Directory structure of the ADF package

Below is the list of major parts of the ADF package.

### 3.1 The bin directory

When the programs have been installed, the binary executables have been placed in \$ADFBIN, typically \$ADFHOMe/bin, with names like 'adf.exe', 'band.exe', and so on. There will also be files 'adf', 'band'. The latter are scripts that execute the binaries.

You should use the script files, rather than the executables directly. The script files take care of several convenient features (like the BASIS key described in the ADF User's Guide) and provide a correct environment for running in parallel. See also the sample run scripts and the Examples document.

The \$ADFBIN/start script takes care of setting up the environment and starting the binaries. If necessary, it parses the information provided by a batch system and sets up a machinefile (or an appfile) and runs tasks in parallel. Edit the file if you need to customize the way MPI programs are started.

### 3.2 The atomicdata directory

The directory atomicdata/ contains a large amount of data needed by programs from the ADF package at run time. For example, it contains basis sets for all atoms. Generally you should **not** modify any of the files in this directory.

The basis sets are described in detail in the ADF BasisSet document, available from the Documentation part of the SCM web site (<http://www.scm.com>).

### 3.3 The examples directory

The directory examples/ contains sample script and output files. The files with extension .run are shell scripts that also contain embedded input. Thus these files should be executed, and not given as input to ADF.

The example calculations are extensively documented in the Examples document, available from the SCM web site (<http://www.scm.com>).

### 3.4 The source code (adf, band, libtc, Install, ...)

The source code files found in the program and library directories and subdirectories thereof have extensions .d, .d90, .dmod\* or .cd and need to be pre-processed by the parser (scu) before compilation. The parser produces files with a suffix .f, .f90 or .c containing plain FORTRAN or C code. These files are placed in the FFILES directory in each program's or library main directory. Other source files (used by the parser) are include files (files with extension .fh or .h) and files that define configuration parameters (the file "settings").

Parsing and compilation are controlled by \$ADFBIN/yam.

The Install directory contains a script configure, some data files which provide generic input for configure (start, starttcl, and some more), a portable preprocessor cpp (based on Mouse cpp), the adf preprocessor (adfparser) and files that will be addressed by the \$ADFBIN/yam command.

Apart from a few more files that are of little importance in the installation, you'll find in `Install/` a number of subdirectories with names like `pentium_linux_ifc`, `sgi_irix65`, and so on: they contain precompiled binaries for supported platforms.

`Install/configure` usually guesses the configuration parameters from the system but in some cases you may want to run it interactively. In this case, start it without command-line parameters and it will ask you some questions: hardware platform, parallel variety (PVM, MPI, serial), and the maximum number of CPU's to use.

Based on the files in the `Install` directory and the chosen configuration parameters, the `configure` script will create in `$ADFHOME/` two files: `settings` and `Makeflags`. It will also populate the `bin` directory with important files.

## 3.5 The Doc directory

All the user documentation for ADF is present in the PDF format in `$ADFBIN/Doc`. The same documentation is also available on the [SCM website](#)

## 3.6 The scripting directory

This directory contains some useful scripts that are part of the ADF package but are distributed under the GNU General Public License.

# 4 Additional Information

## 4.1 Performance-related environment variables

You can customize and optimize your ADF package by changing or setting several environment variables, and/or use some input options. We advise you to test the effects of changing them for some of your typical calculations, and then setting the optimal values in the scripts that define the environment for your login shell. Don't forget to make sure that the changes you make will also apply for batch jobs, and for all nodes of a distributed computing system.

A complete list of environment variables used by the ADF package can be found in the section 2.4. Here we will discuss a few that are particularly important.

### **SCM\_IOBUFFERSIZE**

Most programs within the ADF package use the KF I/O system. This is coupled to a facility to store (parts of) files in memory, if you have enough memory available. Thus, the programs will cache the IO data instead of the operating system, and reduce the amount of IO significantly.

It depends on your operating system and hardware if you can benefit from this scheme. The default buffer size depends on the platform but is typically set to 64, which should be sufficient for running small jobs without much disk I/O. In some cases you can have a major performance improvement by making this buffer much larger, for example 512MB. You do this by setting the SCM\_IOBUFFERSIZE environment variable to the buffer size **in megabytes**.

Please try for yourself, with your typical calculation on your production machine to find out the optimal value for your situation.

### **SCM\_VECTORLENGTH**

Almost all programs within the ADF package use numerical integration, and this is normally the most time-consuming part of the code. Numerical integration involves summing together results for each 'integration point'. By handling a number of points together you can greatly influence the performance. The number of integration points handled together is called the block length.

If the block length is too small, you will have a significant overhead and the programs will be very slow.

If the block length is too large, lots of data will not fit in cache memory and again the program will not run at the maximum speed.

The optimal block length is somewhere in between, ranging from 32 to 4096 depending on your hardware. Sometimes it pays off to set the block length explicitly NOT to a power of 2 to avoid memory bank conflicts.

Again, try for yourself, with your typical calculation on your production machine to find out the optimal value for your situation. On most machines, 128 is a good compromise value.

### **SCM\_TMPDIR**

This environment variable determines where all processes create their scratch directories. Setting this variable correctly is extremely important for performance reasons. If this variable is not set, the scratch directories will be created in the current working directory. For a detailed description see section 2.8.

## 4.2 MPI

MPI (Message Passing Interface) is a standard describing how to pass messages between programs running on the same or different machines.

MPI is a formal standard and it is actively supported by all major vendors. Some vendors have highly-optimized MPI libraries available on their systems. There are also a couple of open-source implementations of the MPI standard, such as MPICH and OpenMPI. There are also numerous commercial MPI implementations that support a wide range of systems and interconnects, for example, HP-MPI and IntelMPI.

Support for a particular MPI implementation in ADF can be considered at three levels: the source code, the configure script, and pre-compiled binaries. At each level different MPI implementations may be supported.

The ADF source code is not implementation-specific and thus theoretically it supports any MPI library. Many popular MPI implementations are supported at the level of the configure script. For example on 32-bit Linux these are: MPICH1, Intel MPI, OpenMPI, HP-MPI, LAM-MPI, and Score. This means that proper compiler flags will be used and an appropriate \$ADFBIN/start script will be generated at configure time.

However, when choosing an MPI implementation for pre-compiled binaries SCM considers many factors including (but not limited to) the re-distribution policy, performance, and built-in support for modern interconnects. HP-MPI is currently the standard MPI implementation supported by SCM because it has the most favorable combination of these factors at this moment. For platforms where HP-MPI is supported it is distributed with ADF. A different MPI implementation will be standard on a platform where HP-MPI is not available. It may or may not be distributed with ADF. For example, SGI MPT is standard on SGI machines and OpenMPI is standard on MacOS platforms, but only the latter is distributed together with ADF.

When pre-compiled binaries do not work on your computer(s) due to incompatibility of the standard MPI library with your soft- and/or hardware, the SCM support staff will be glad to assist you in compiling ADF with the MPI implementation supported on your machine(s).

If you are going to use an MPI version of the ADF package, and it not HP-MPI or OpenMPI, you will need to determine if the corresponding MPI run-time environment is already installed on your machine. If not, you will need to install it separately from ADF. As it has been already mentioned, HP-MPI and OpenMPI are distributed together with the corresponding version of ADF so you don't need to worry about installing them separately.

### Running with MPI on more than one node

When running on more than one machine (for example on a cluster **without** a batch system) you need to specify a list of hosts on which mpirun needs to spawn processes. In principle, this is implementation-specific and may be not required if the MPI is tightly integrated with your operating and/or batch system. For MPICH1 and HP-MPI, you can do this by preparing a file containing hostnames of the nodes (one per line) you will use in your parallel job. Then you set the SCM\_MACHINEFILE environment variable pointing to the file.

When you submit a parallel job to a batch system the job scheduler usually provides a list of nodes allocated to the job. The \$ADFBIN/start shell script has some logic to extract this information from the batch system and pass it to the MPI's launcher command (typically mpirun). In some cases, depending on your cluster configuration, this logic may fail. If this happens, you should examine the \$ADFBIN/start file and edit the relevant portion of it. For example, you may need to change commands that process the batch system-provided nodelist or change mpirun's command-line options or even replace the mpirun command altogether.

## 4.3 More on HP-MPI

As it has been already mentioned, HP-MPI is currently the standard MPI implementation on some systems, namely under Windows and on Intel-based Linux machines. Full documentation about HP-MPI is available at the [HP website](#). HP-MPI is distributed with (HP-MPI versions of) ADF and a complete HP-MPI directory tree is found in \$ADFBIN/hpmpi. Using HP-MPI with ADF does not require any additional licenses. In addition to TCP/IP and shared memory supported by every MPI implementation, HP-MPI also supports, without recompilation, all well-known interconnects such as Infiniband and Myrinet. The best available interconnect is chosen automatically at run time and this can also be modified using mpirun command-line switches in the \$ADFBIN/start script.

A few words about the mpirun commands found in the start script. If you look inside the \$ADFBIN/start file on Linux, you will likely see three mpirun commands. They differ in the way the list of nodes is specified, if specified at all. The mpirun command with a `-lsb_hosts` switch is used under the LSF batch system. It employs the HP-MPI's tight integration with LSF and lets it pick up the job configuration directly from the batch system.

The mpirun command with a `-hostfile` switch is used for multi-node interactive jobs and for all parallel jobs started under PBS or SGE batch systems. In the case of a batch system, the hostfile provided by the scheduler is passed to the mpirun command as-is, that is without modifications. The latter can be a problem if the format of the hostfile is different from that supported by HP-MPI, in which case it should be modified accordingly. Please see the HP-MPI User's Guide for information about supported hostfile formats.

Finally, the simplest form of the mpirun command is used for single-node jobs when the calculation is performed on the localhost. In this case the NSCM environment variable determines how many parallel processes will be started. Please note that if "\$NSCM" is not equal to "1" (digit one) then the exact value of the NSCM variable has no effect when the job is started under any of the batch systems or with a hostfile. In this case the number of parallel processes to start is determined by the contents of the hostfile.

One important point about HP-MPI is that it needs a remote shell command, such as ssh, to start tasks on compute nodes. By default ssh is used but this can be changed using the MPI\_REMSH environment variable. For example, executing the following from the bash prompt will make mpirun use rsh instead of ssh:

```
export MPI_REMSH=rsh
```

As usual, if you want to make the change permanent, you need to add this command to a shell resource file.

There are cases when a Linux cluster is configured in such a way that both ssh and rsh communication from/to/between compute nodes is disabled. One of the most common examples is TORQUE with the MAUI scheduler. In this case, there is a remote shell replacement utility called *pbsdsh*. This utility checks that it is executed under PBS and, if yes, allows you to start remote programs only on the nodes allocated to that particular job. In principle, this is all that mpirun needs. The only problem is that *pbsdsh* uses different command-line options. To solve this, we provide in \$ADFBIN a script called *torque\_ssh*. To make HP-MPI always use *torque\_ssh* instead of ssh simply set MPI\_REMSH in the shell resource file as follows (of course **after** the ADFBIN definition):

```
MPI_REMSH=$ADFBIN/torque_ssh
export MPI_REMSH
```

## 4.4 PVM

**NOTE:** SCM strongly discourages the use of PVM. It is not actively developed, it does not use recent fast interconnects and it is significantly slower than any modern MPI implementation. So it should only be used as the last resort when no other parallel option works.

PVM (Parallel Virtual Machine) provides a common interface for a large number of different parallel architectures, ranging from a heterogeneous cluster of workstations to shared memory multiprocessors. It is public software, and you can download it for free. Some vendors have their own optimized implementations of PVM to exploit specific hardware features. Unfortunately, these special versions are not always flawless.

If you are not familiar with PVM and need to install it yourself, we strongly suggest you have a look at the PVM documentation. PVM (including FAQ, User's Guide and the Tutorial) can be downloaded from:

<http://www.netlib.org/pvm3/index.html>

After you've installed PVM make sure the PVM daemon pvmd can be started on all nodes of your system. If pvmd fails to start on (some) nodes you may have to look at .rhosts or ssh/rsh issues. Also make sure that there is no firewall blocking communication between nodes. In case of trouble, you can also find some error messages in the /tmp/pvml.\$UID file on the failing node.

To run in parallel you start the programs as usual, but you (or your batch system, job script, ...) need to start and stop the PVM virtual machine yourself. It is convenient to make a special 'submit' script that handles this for you if you are using a batch system.

## 4.5 Known HP-MPI/Linux issue: libibverbs error message

HP-MPI 2.2.7 stops with the following (or similar) error message

libibverbs: Fatal: couldn't open sysfs class 'infiniband\_verbs'.

This occurs, for example, on GigE-connected ROCKS clusters with OpenMPI installed. The error is caused by the fact that there are Infiniband libraries installed without corresponding kernel drivers and/or hardware. In this case, one has to enforce the use of TCP by HP-MPI, which can be done using one of the methods below:

- edit the \$ADFBIN/start script and add a -tcp option to all mpirun commands:  
\$ADFBIN/hpmi/bin/mpirun.mpich -TCP ...
- or set the MPIRUN\_OPTIONS environment variable in your shell resource file (for example, ~/.bash\_profile):  
export MPIRUN\_OPTIONS=-TCP
- or set the MPI\_IC\_ORDER environment variable in your shell resource file (for example, ~/.bash\_profile):  
export MPI\_IC\_ORDER="TCP"

Any of these options will make sure all other interconnects but TCP are ignored.