# Intel® Fortran Compiler for Linux* Systems User's Guide
# Volume I: Building Applications

# Disclaimer and Legal Information

Information in this document is provided in connection with Intel® products. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted by this document. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. Intel products are not intended for use in medical, life saving, or life sustaining applications.

This *Intel® Fortran Compiler for Linux\* Systems User's Guide Volume I: Building Applications* as well as the software described in it is furnished under license and may only be used or copied in accordance with the terms of the license. The information in this manual is furnished for informational use only, is subject to change without notice, and should not be construed as a commitment by Intel Corporation. Intel Corporation assumes no responsibility or liability for any errors or inaccuracies that may appear in this document or any software that may be provided in association with this document.

Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

The Intel Fortran Compiler product may contain software defects which may cause the product to deviate from published specifications. Current characterized software defects are available on request.

Intel SpeedStep, Intel Thread Checker, Celeron, Dialogic, i386, i486, iCOMP, Intel, Intel logo, Intel386, Intel486, Intel740, IntelDX2, IntelDX4, IntelSX2, Intel Inside, Intel Inside logo, Intel NetBurst, Intel NetStructure, Intel Xeon, Intel XScale, Itanium, MMX, MMX logo, Pentium, Pentium II Xeon, Pentium III Xeon, Pentium M, and VTune are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

\* Other names and brands may be claimed as the property of others.

Copyright © Intel Corporation 2003-2004.
Portions © Copyright 2001 Hewlett-Packard Development Company, L.P.

# Table Of Contents

Table Of Contents

Table Of Contents

Table Of Contents

Table Of Contents

# Table Of Contents

Table Of Contents

Table Of Contents

Table Of Contents

Table Of Contents

Table Of Contents

xx

Table Of Contents

# About the Intel® Fortran Compiler

The Intel® Fortran Compiler version 8.0 compiles code targeted for the IA-32 Intel® architecture and Intel® Itanium® architecture.

The Intel Fortran Compiler product includes the following components for the development environment:

- Intel Fortran Compiler for 32-bit Applications
- Intel Fortran Compiler for Itanium-based Applications
- Intel Debugger (idb)

The Intel Fortran Compiler for Itanium-based applications includes the Intel Itanium Assembler and Intel Itanium Linker.

See also How to Use This Document.

# How to Use This Document

This is Volume I in the two-volume *Intel® Fortran Compiler for Linux\* Systems User's Guide*. It explains how you can use the Intel Fortran Compiler to build applications. Volume II explains how to optimize applications.

This User's Guide provides information on how to get started with Intel Fortran, how the compiler operates, and how to develop applications.

This documentation assumes that you are familiar with the Fortran Standard programming language and with the Intel® processor architecture. You should also be familiar with the host computer's operating system.

> **Note**
> This document explains how information and instructions apply differently to each targeted architecture. If there is no specific indication to either architecture, the description is applicable for both architectures.

## Additional Documentation

Besides the two volumes of the *User's Guide*, you should also have access to these manuals:

- *Intel®Fortran Compiler Options Quick Reference Guide for Linux\* Systems*
- *Intel® Fortran Language Reference*

Intel® Fortran Compiler for Linux* Systems User's Guide Volume I: Building Applications

- *Intel® Fortran Libraries Reference*
- *Intel® Fortran Release Notes*

# Notation Conventions

This manual uses the following conventions.

| | |
|---|---|
| Intel® Fortran | The name of the common compiler language supported by the Intel® Visual Fortran Compiler for Windows* and Intel Fortran Compiler for Linux* products. |
| `This type style` | Elements of syntax, reserved words, option keywords, variables, file names, and code examples are shown in a monospaced font. The text appears in lowercase unless uppercase is required. |
| THIS TYPE STYLE | Statements, keywords, and directives are shown in all uppercase, in a normal font. For example, "add the USE statement…" |
| **This type style** | Bold normal text shows menu names, menu items, button names, dialog window names, and other user-interface items. |
| **File>Open** | Menu names and menu items joined by a greater than (>) sign indicate a sequence of actions. For example, "Click **File>Open**" indicates that in the **File** menu, click **Open** to perform this action. |
| **`This type style`** | Bold, monospaced text indicates user input. Shows what you type as command or input. |
| *`This type style`* | Italic, monospaced text indicates placeholders for information that you must supply. Italics are also used to introduce new terms. |
| `[options]` | Items inside single square brackets are optional. (In some examples, square brackets are used to show arrays.) |
| `{value | value}` | Braces and a vertical bar indicate a choice among two or more items. You must choose one of the items unless all of the items are also enclosed in square brackets. |
| ... | A horizontal ellipsis (three dots) following an item indicates that the item preceding the ellipsis can be repeated. In code examples, a horizontal ellipsis means that not all of the statements are shown. |
| Microsoft* Windows XP* | An asterisk at the end of a word or name |

| | indicates it is a third-party product trademark. |
|---|---|

## Getting Started

# Getting Started Overview

See these topics:

How You Can Use the Intel Fortran Compiler

Compilation Phases

Preprocess Phase

Assemblers and Linker

Default Behavior of the Intel Fortran Compiler

Input Files and Filename Extensions

File Specifications

Output Files

Temporary Files Created by the Compiler or Linker

# How You Can Use the Intel® Fortran Compiler

The Intel® Fortran Compiler has the following variations:

- The Intel® Fortran Compiler for 32-bit Applications is designed for IA-32 systems. The IA-32 compilations run on any IA-32 Intel processor and produce applications that run on IA-32 systems. This compiler can be optimized specifically for one or more Intel® IA-32 processors, such as Pentium M, Pentium 4, and Xeon™.
- The Intel® Fortran Itanium® Compiler for Itanium® -based Applications, or native compiler, is designed for Itanium architecture systems. This compiler runs on Itanium-based systems and produces Itanium-based applications. Itanium-based compilations can only operate on Itanium-based systems.

The command to invoke either of these compilers is `ifort`.

The Intel® Fortran Compiler has a variety of options that enable you to use the compiler features for higher performance of your application.

The Intel® Fortran Compiler enables your software to perform the best on Intel architecture-based computers. The compiler has several high-performance optimizations. Some of its features and benefits are:

| What feature might you want to use? | How will this help you? |
| --- | --- |
| Support for Streaming SIMD Extensions (SSE) and Streaming SIMD Extensions 2 (SSE2) | Intel microarchitecture benefit. |
| Automatic vectorizer | Parallelism in your code achieved automatically. |
| Parallelization | Automatic generation of multithreaded code for loops. Shared memory parallel programming with OpenMP*. |
| Floating-point optimizations | Improved floating-point performance. |
| Data prefetching | Improved performance due to the accelerated data delivery. |
| Interprocedural optimizations | Better performance for larger applications. |
| Whole program optimization | Improved performance between modules in larger applications. |
| Profile-guided optimization | Improved performance based on profiling the frequently used procedures. |
| Processor dispatch | Use of the latest Intel architecture features while maintaining object code compatibility with previous generations of Intel® Pentium® processors. |

# Compilation Phases

The compiler processes Fortran language source and generates object files. You decide the input and output by setting options when you run the compiler.

When invoked, the compiler determines which compilation phases to perform based on the extension of the source filename and on the compilation options specified in the command line.

The compilation phases and the software that controls each phase are shown below:

| Compilation phase | Controlling software | IA-32 or Itanium®-based applications? |
|---|---|---|
| Preprocess (optional) | `fpp` | Both |
| Compile | `fortcom` | Both |
| Assemble (optional) | `as` or `ias` | `as` for IA-32 applications; `ias` for Itanium-based applications |
| Link | `ld(1)` | Both |

By default, the compiler generates an object file directly without calling the assembler. However, if you need to use specific assembly input files and then link them with the rest of your project, you can use an assembler for these files.

The compiler passes object files and any unrecognized filename to the linker. The linker then determines whether the file is an object file (`.o`) or a library (`.a`) or shared library (`.so`).The compiler handles all types of input files correctly. Thus, you can use it to invoke any phase of compilation.

# Preprocess Phase

Preprocessing performs such tasks as preprocessor symbol (macro) substitution, conditional compilation, and file inclusion. The compiler preprocesses files as an optional first phase of the compilation. Source files that have a filename extension of `.fpp`, `.F`, `.F90`, `.FOR`, `.FTN`, or `.FPP` are automatically preprocessed by the compiler. For example, the following command preprocesses a source file that contains standard Fortran preprocessor directives, then passes the preprocessed file to the compiler and linker:

**`ifort source.fpp`**

If you want to preprocess files that have another extension, you have to overtly specify the preprocessor.

You do not usually need to specify preprocessing for Fortran source programs. The preprocessor is necessary only if your program uses C-style preprocessing commands, such as `#if`, `#define`, and so forth.

If you choose to preprocess your source programs, you must use the preprocessor fpp, which is the preprocessor supplied with the Intel® Fortran Compiler, or the preprocessing capability of a Fortran compiler. It is recommended that you use fpp.

If you want to use another preprocessor, you must invoke it before you invoke the compiler.

fpp conforms to cpp and accepts the cpp-style directives. cpp (and thus fpp) prohibit the use of a string constant value in an `#if` expression.

You can use the Preprocessor Options on the command line to direct the operations of the preprocessor.

### ⚠️Caution

Using a preprocessor that does not support Fortran can damage your Fortran code, especially with FORMAT statements. For example, `FORMAT (\\I4)` changes the meaning of the program because the backslash "\" indicates end-of-record.

# Assemblers and Linker

The assemblers and linker you can use are summarized in this table:

| Tool | Default | Provided with Intel Fortran Compiler? |
|------|---------|---------------------------------------|
| IA-32 assembler | Linux* assembler, `as` | No |
| Itanium® assembler | Intel® Itanium® assembler, `ias` | Yes |
| Linker | System linker, `ld(1)` | No |

You can specify alternate tool locations and options for preprocessing, compilation, assembly, and linking.

See also Libraries Provided by Intel Fortran.

## Assemblers

For 32-bit applications, Linux supplies its own assembler, `as`.

For Itanium-based applications, use the Itanium assembler, `ias`. For example, to link some specific input file to the Fortran project object file, do the following:

1. Issue a command using the -S option to generate an assembly code file, `file.s`: **ifort -S -c *file.f***
2. To assemble the `file.s` file, call the Itanium® assembler with this command: **ias -Nso -p32 -o *file.o file.s***

where the following assembler options are used:

`-Nso` suppresses the sign-on message.

`-p32` enables defining 32-bit elements as relocatable data elements. (This option is available for backward compatibility.)

`-o file.o` indicates the output object filename.

## Linker

The compiler calls the system linker, `ld(1)`, to produce an executable file from the object files.

# Default Behavior of the Intel Fortran Compiler

The compiler generates one or more executable files of one or more input files. By default, it performs the following actions:

- Searches for all files, including library files, in the current directory.
- Passes options designated for linking to the linker.
- Passes user-defined libraries to the linker.
- Displays error and warning messages.
- Performs default settings and optimizations, unless these options are overridden by specific options settings.
- For IA-32 applications, uses the `-tpp7` option to optimize the code for the Intel® Pentium® 4 and Intel® Xeon™ processor.
- For Itanium® -based applications, uses the `-tpp2` option to optimize the code for the Intel® Itanium® 2 processor.

### Note

On operating systems that support characters in Unicode* (multi-byte) format, the compiler will process file names containing Unicode* characters.

# Input Files and Filename Extensions

The Intel Fortran Compiler interprets the type of each input file by its filename extension, such as `.a`, `.f`, `.for`, `.o`, and so on:

| Filename | Interpretation | Action |
|---|---|---|
| `filename.a` | Object library | Passed to `ld`. |
| `filename.f`<br>`filename.ftn`<br>`filename.for`<br>`filename.i` | Fortran fixed-form source | Compiled by the Intel® Fortran compiler. |
| `filename.fpp`<br>`filename.F`<br>`filename.F90`<br>`filename.FOR`<br>`filename.FTN`<br>`filename.FPP` | Fortran fixed-form source | Preprocessed by the Intel Fortran preprocessor `fpp`; then compiled by the Intel Fortran compiler. |
| `filename.f90`<br>`filename.i90` | Fortran free-form source | Compiled by the Intel Fortran compiler. |
| `filename.s` | Assembly file | Passed to the assembler (IA-32 compiler) or the Intel Itanium® assembler (Itanium-based compiler). |
| `filename.o` | Compiled object file | Passed to `ld`. |

You can use the compiler configuration file to specify default directories for input libraries. To specify additional directories for input files, temporary files, libraries, and for the files used by the assembler and the linker, use compiler options that specify output file and directory names.

# File Specifications

A complete file specification consists of a file name usually preceded by a pathname that specifies a directory. The pathname can be in one of two forms:

- An absolute pathname, where the directory is specified relative to the root directory. The first character is a slash (`/`). For example, the following directory and file name refer to the file named `testdata` in the `/usr/users/gdata` directory: `/usr/users/gdata/testdata`
- A relative pathname, where the specified directory is relative to the current directory. Relative pathnames do not begin with a slash (`/`). The following example uses a relative pathname from the current directory `/usr/users` to refer to the same file `testdata` in the `gdata/` subdirectory: `gdata/testdata`

Directory names and file names should not contain any operating system wildcard characters (such as `*`, `?`, and the `[ ]` construct). You can use the tilde (~) character as the first character in a pathname to refer to a top-level directory as in the C shell.

When specifying files, keep in mind that trailing and leading blanks are removed from character expression names, but not from Hollerith (numeric array) names.

File names are case-sensitive and can consist of uppercase and lowercase letters. For example, the following file names represent three different files:

```
myfile.for
MYfile.for
MYFILE.for
```

# Output Files

The output produced by the `ifort` command includes:

- An object file (such as `test.o`), if you specify the `-c` option on the command line. An object file is created for each source file.
- An executable file (such as `a.out`), if you omit the `-c` option.
- One or more module files (such as `datadef.mod`), if the source file contains one or more MODULE statements.

You control the production of these files by specifying the appropriate options on the command line.

The compiler generates a temporary object file for each source file, unless you specify the `-c` option. The linker is then invoked to link the object files into one executable program file and the temporary object files are deleted.

If you specify the `-c` option, the object files are created and retained in the current working directory. You must link the object files later by using a separate `ifort` command. This allows incremental compilation of a large application, perhaps by means of a makefile processed by the `make` command.

If fatal errors are encountered during compilation, or if you specify certain options such as `-c`, linking does not occur.

## Note

To compile all objects over the entire program, use the `-ipo` option.

To specify a file name for the executable program file (other than `a.out`), use the `-o output` option, where *output* specifies the file name. The following command requests a file name of `prog1.out` for the source file `test1.f`:

```
ifort -o prog1.out test1.f
```

If you specify the `-c` option with the `-o` *output* option, you rename the object file (not the executable program file). If you specify `-c` and omit the `-o` *output* option, the compiler names the object files with a `.o` suffix substituted for the source file suffix.

### Note

You cannot use `-c` and `-o` together with multiple source files.

The default optimization level is `-O2` (unless you specify `-g`).

# Temporary Files Created by the Compiler or Linker

Temporary files created by the compiler or linker reside in the directory used by the operating system to store temporary files.

To store temporary files, the driver first checks for the `TMP` environment variable. If defined, the directory that `TMP` points to is used to store temporary files.

If the `TMP` environment variable is not defined, the driver then checks for the `TMPDIR` environment variable. If defined, the directory that `TMPDIR` points to is used to store temporary files.

If the `TMPDIR` environment variable is not defined, the driver then checks for the `TEMP` environment variable. If defined, the directory that `TEMP` points to is used to store temporary files.

If the `TEMP` environment variable is not defined, the `/tmp` directory is used to store temporary files.

## Building Applications

# Building Applications Overview

See these topics about Intel® Fortran:

Controlling the Compilation Process

Setting and Viewing Environment Variables

# Controlling the Compilation Process

To customize the environment used during compilation, you can specify variables, options, and files as follows:

- Environment variables to specify paths where the compiler searches for special files such as libraries and "include" files
- Configuration files to specify the options used for every compilation and response files to specify the options and files used for individual projects

# Setting and Viewing Environment Variables

You can use the SET command to view or set environment variables one at a time. You can also set environment variables by using the `ifortvars.csh` and `ifortvars.sh` files to set several at a time. The files are found in this

directory: `/opt/intel_fc_80/bin`. See Running the Shell Script to Set Up the Environment Variables.

Within the C Shell, use the `setenv` command to set an environment variable:

**setenv FORT8 /usr/users/smith/test.dat**

To remove the association of an environment variable and its value within the C shell, use the `unsetenv` command.

**unsetenv FORT8**

Within the Bourne shell (`sh`), the Korn shell (`ksh`), and the bash shell, use the `export` command and assignment command to set the environment variable:

**export FORT8**
**FORT8=/usr/users/smith/test.dat**

To remove the association of an environment variable and its value within the Bourne shell, the Korn shell, or the bash shell, use the `unset` command:

**unset FORT8**

## Configuration File Environment Variables

By default, the compiler picks up the default configuration file (`ifort.cfg`) from the same directory where the compiler executable resides. However, if you want the compiler to use another configuration file in a different location, you can use the `IFORTCFG` environment variable to assign the directory and filename for the other configuration file.

**See Also**

Compile-Time Environment Variables

Run-Time Environment Variables

# Running the Shell Script to Set Up the Environment Variables

Before you first invoke the compiler, you need to set the environment variables to specify locations for the various components.

The Intel Fortran Compiler installation includes a shell script that you can use to
set environment variables.

Use the `source` command to execute the shell script from the command line.
For example, to execute this script file for the bash shell:

**source /opt/intel_fc_80/bin/ifortvars.sh**

If you use the C shell, use the `.csh` version of this script file:

**source /opt/intel_fc_80/bin/ifortvars.csh**

If you want `ifortvars.sh` to run automatically when you start Linux*, edit your
`.bash_profile` file and add the line above to the end of your file. For example:

```
# set up environment for Intel compiler
source /opt/intel_fc_80/bin/ifortvars.sh
```

# Invoking the Intel Fortran Compiler

You can invoke the Intel® Fortran Compiler in either of two ways:

- Using the `ifort` command
- Using the `make` command to specify a makefile

## Using the ifort Command

The syntax is:

**ifort [*options*] *input_file(s)***

An *option* is specified by one or more letters preceded by a hyphen.

Some options take arguments in the form of filenames, strings, letters, or
numbers. Except where otherwise noted, you can enter a space between the
option and its argument(s) or you can combine them. See Compiler Options
Overview.

You can specify more than one *input_file*, using a space as a delimiter. See
Input Files and Filename Extensions.

**Note**

Options on the command line apply to all files. For example, in the following command line, the `-c` and `-nowarn` options apply to both files `x.f` and `y.f`:
```
ifort -c x.f -nowarn y.f
```

## Using the make Command

To compile a number of files with various paths and to save this information for multiple compilations, you can use a makefile to invoke the Intel® Fortran Compiler.

To use a makefile to compile your input files, make sure that `/usr/bin` and `/usr/local/bin` are in your path.

If you use the C shell, you can edit your `.cshrc` file and add the following:

```
setenv PATH  /usr/bin:/usr/local/bin:yourpath
```

Then you can compile as:

**make -f *yourmakefile***

where `-f` is the `make` command option to specify a particular makefile.

# Examples of the ifort Command

## Compiling and Linking Multiple Files

The following `ifort` command compiles the Fortran free-format source files `aaa.f90`, `bbb.f90`, and `ccc.f90`. The command invokes the `ld` linker and passes the temporary object file to the linker, which it uses to produce the executable file `a.out`:

**ifort aaa.f90 bbb.f90 ccc.f90**

The following `ifort` command compiles all file names that end with `.f.` as Fortran fixed-format source. The linker produces the `a.out` file:

**ifort *.f**

## Preventing Linking

The following `ifort` command compiles, but does not link, the free-format source file `typedefs_1.f90`, which contains a MODULE TYPEDEFS_1. The command creates files `typedefs_1.mod` and `typedefs_1.o`. The object file is retained automatically. Specifying the `-c` option prevents linking:

```
ifort -c typedefs_1.f90
```

# Compiling Fortran 95/90 and C Source Files

The following `ifort` command compiles the free-format Fortran main program `myprog.f90`., which references the module `TYPEDEFS_1` with a USE TYPEDEFS_1 statement (it uses the object file created in the previous example). The module file `typedefs_1.mod` is read from the current directory. The main program calls a function written in C. The C routine named `utilityx_` is declared in a file named `utilityx.c`:

```
ifort myprog.f90 typedefs_1.o utilityx.c
```

`ifort` does not recognize a source file with a `.c` extension as needing to be compiled by the C compiler. Instead, it passes it as an "unknown" to the linker. You will need to call the C compiler to compile `utilityx.c`.

# Renaming the Output File

The following `ifort` command compiles the free-format Fortran source files `circle-calc.f90` and `sub.f90` together:

```
ifort -c circle-calc.f90 sub.f90
```

The default optimization level `-O2` applies to both source files during compilation. Because the `-c` option is specified, the object files are not passed to the linker. In this case, the named output files are the object files.

Like the previous command, the following `ifort` command compiles multiple source files:

```
ifort -o circle.out circle-calc.f90 sub.f90
```

Because the `-c` option was omitted, an executable program named `circle.out` is created.

# Specifying an Additional Linker Library

The following `ifort` command compiles a free-format source file `myprog.f90` using default optimization, and passes an additional library for the linker to search:

**ifort typedefs_1.o myprog.f90 -lmylib**

The file is processed at optimization level `-O2` and then linked with the object file typedefs_1.o . The `-lmylib` option instructs the linker to search in the `libmylib` library for unresolved references (in addition to the standard list of libraries the `ifort` command passes to the linker).

# Using Module (.mod) Files

A module (`.mod` file) is a type of program unit that contains specifications of such entities as data objects, parameters, structures, procedures, and operators. These precompiled specifications and definitions can be used by one or more program units. Partial or complete access to the module entities is provided by the USE statement. Typical applications of modules are the specification of global data or the specification of a derived type and its associated operations.

Some programs require modules located in multiple directories. You can use the `-Idir` option when you compile the program to locate the `.mod` files that should be included in the program.

You can use the `-module path` option to specify the directory where to create the module files. This path is also used to locate module files. If you don't use this option, module files are created in the default path.

You need to make sure that the module files are created before they are referenced by another program or subprogram.

## Compiling Programs with Modules

If a file being compiled has one or more modules defined in it, the compiler generates one or more `.mod` files. For example, a file `a.f90` contains modules defined as follows:

```
module test
integer:: a
contains
subroutine f()
end subroutine
end module
module payroll
.
```

```
.
.
end module
```

This compiler command:

**ifort -c a.f90**

generates the following files:

- `test.mod`
- `test.o`
- `payroll.mod`
- `payroll.o`

The `.mod` files contain the necessary information regarding the modules that have been defined in the program `a.f90`.

If the program does not contain a module, no `.mod` file is generated. For example, `test2.f90` does not contain any modules. This compiler command:

**ifort -c test2.f90**

produces just an object file, `test2.o`.

For another example, assume that `file1.f90` contains one or more modules and `file2.f90` contains one or more program units that access these modules with the USE statement. The sources can be compiled and linked by this command:

**ifort file1.f90 file2.f90**

## Working with Multi-Directory Module Files

For an example of managing modules when the `.mod` files could be produced in different directories, assume that the program `mod_def.f90` resides in directory `/usr/yourdir/test/t`, and this program contains a module defined as follows:

```
file: mod_def.f90
module definedmod
.
.
.
end module
```

The compiler command:

```
ifort -c mod_def.f90
```

produces two files: `mod_def.o` and `definedmod.mod` in directory `/usr/yourdir/test/t`.

If you need to use the above `.mod` file in another directory, for example, in directory `/usr/yourdir/test/t2`, where the program `usemod` uses the `definedmod.mod` file, do the following:

```
file: use_mod_def.f90
program usemod
use definedmod
.
.
.
end program
```

To compile the above program, use this command:

```
ifort -c use_mod_def.f90 -I/usr/yourdir/test/t
```

where the `-Idir` option provides the compiler with the path to search and locate the `definedmod.mod` file.

# Parallel Invocation with a makefile

The programs in which modules are defined support the compilation mechanism of parallel invocation with a makefile for interprocedural optimizations of multiple files and of the whole program. Consider the following code:

```
test1.f90
module m1
.
.
.
end module
test2.f90
subroutine s2()
use m1
.
.
.
end subroutine
test3.f90
subroutine s3()
use m1
.
```

```
.
.
end subroutine
```

The makefile to compile the above code looks like this:

```
m1.mod: test1.o
test1.o:
ifort -c test1.f90
test2.o: m1.mod
ifort -c test2.f90
test3.o: m1.mod
ifort -c test3.f90
```

# Searching for Include and .mod Files

Include files are brought into a program with the `#include` preprocessor directive or a Fortran INCLUDE statement.

Directories are searched for include files in this order:

1. Directory of the source file that contains the include
2. Directories specified by the `-Idir` options
3. Current working directory
4. Directories specified with the `FPATH` environment variable

The locations of directories to be searched are known as the include file path. More than one directory can be specified in the include file path.

A module (`.mod`) file is specified in a program by a USE statement. Module files can be located in multiple directories.

Directories are searched for `.mod` files in this order:

1. Directory of the source file that contains the USE statement
2. Directories specified by the `-module path` option
3. Directories specified by the `-Idir` option
4. Current working directory
5. Directories specified with the `FPATH` environment variable

## Specifying and Removing an Include File Path

You can use the `-Idir` option to indicate the location of include files and module files.

To prevent the compiler from searching the default path specified by the `FPATH` environment variable, use the `-X` option.

You can specify these options in the configuration file, `ifort.cfg`, or on the command line.

For example, to direct the compiler to search the path `/alt/include` instead of the default path, use the following command line:

**-X -I /alt/include newmain.f**

# Configuration Files and Response Files

*Configuration files* and *response files* are slightly different variations in the same idea--the concept that you can use files with various options to eliminate the need to enter the same commands again and again. (Response files are also known as indirect command files.)

## Configuration Files

You can use a configuration (`.cfg`) file to:

- Decrease the time you spend entering command-line options
- Ensure consistency of often used commands

You can insert any valid command-line options into a configuration file. The compiler processes options in the configuration file in the order in which they appear followed by the command-line options that you specify when you invoke the compiler.

### 📝Note

Options placed in the configuration file will be included each time you run the compiler. If you have varying option requirements for different projects, use response files.

By default, a configuration file named `ifort.cfg` is used.

This file resides in the same directory where the compiler executable resides.

However, if you want the compiler to use another configuration file in a different location, you can use the `IFORTCFG` environment variable to assign the directory and file name for the other configuration file.

## Example Configuration File

An example configuration file is shown below. The pound (#) character indicates that the rest of the line is a comment.

```
## Example ifort.cfg file
##
## Define preprocessor macro MY_PROJECT.
-DMY_PROJECT
##
## Set extended-length source lines.
-extend_source
##
## Set maximum floating-point significand precision.
-pc80
##
```

# Response Files

You can use response files (also known as indirect command files) to:

- Specify options used during particular compilations for particular projects
- Save this information in individual files

Response files are invoked as an option on the command line. Options specified in a response file are inserted in the command line at the point where the response file is invoked.

Like configuration files, response files are used to:

- Decrease the time you spend entering command-line options
- Ensure consistency of often used commands

However, options in a configuration file are executed every time you run the compiler. In contrast, you use response files to maintain options for individual projects.

You can place any number of options or file names on a line in the indirect command file or response file. Several files can be referenced in the same command line.

The syntax for using response files is:

**`ifort @responsefile [@responsefile2...]`**

> 📒**Note**

> An "at" sign (@) must precede the name of the response file on the command line.

# Specifying Alternative Tool Locations and Options

The Intel® Fortran compiler lets you specify alternative tool locations and tool options to be used instead of default tools for preprocessing, compilation, assembly, and linking. You can use command-line options to do this.

## Using -Qlocation to Specify an Alternative Location for a Tool

`-Qlocation` lets you specify the pathname location of a supporting preprocessor, compiler, assembler, or linker. This option's syntax is:

`-Qlocation,tool,path`

where `tool` is:

- `fpp`     Intel Fortran preprocessor (fpp)
- `f`       Intel Fortran compiler (fortcom)
- `c`       Intel C compiler
- `as`     Assembler
- `link`    Linker

and `path` is the location of the tool.

**Example**:

**`ifort -Qlocation,fpp,/usr/preproc myprog.f`**

## Using -Qoption to Pass Options to Tools

`-Qoption` lets you pass options to the preprocessor, compiler, assembler, or linker. This option's syntax is:

`-Qoption,tool,options`

where `tool` is:

- `fpp`    Intel Fortran preprocessor (fpp)
- `f`      Intel Fortran compiler (fortcom)
- `c`      Intel C compiler
- `as`     Assembler
- `link`   Linker

and `options` is one or more valid argument strings for the designated tool.

If the argument contains a space or tab character, you must enclose the entire argument in quotation marks (" "). You must separate multiple arguments with commas.

The following example directs the linker to link with an alternative library:

```
ifort -Qoption,link,-lmylib prog1.f
```

# Predefined Preprocessor Symbols

Preprocessor symbols (macros) let you substitute values in a program before it is compiled. The substitution is performed in the preprocessing phase.

Some preprocessor symbols are predefined by the compiler system and are available to compiler directives and fpp. If you want to use others, you need to specify them on the command line.

See also Preprocessor Options.

The predefined preprocessor symbols available for the Intel® Fortran compiler are described in the table below. The Default column describes whether the preprocessor symbol is enabled (on) or disabled (off) by default.

| Symbol Name | Default | Architecture | Description |
|---|---|---|---|
| `__INTEL_COMPILER=`$n$ | On, $n$=800 | Both | Identifies the Intel Fortran Compiler |
| `__linux__`<br>`__linux`<br>`__gnu_linux__`<br>`linux`<br>`__unix__`<br>`__unix`<br>`unix`<br>`__ELF__` | | Both | Defined at the start of compilation |
| `__i386__` | | IA-32 | |

| | | | |
|---|---|---|---|
| __i386 | | | |
| i386 | | | |
| __ia64__ | | Itanium® | |
| __ia64 | | | |
| ia64 | | | |
| `_OPENMP=n` | *n*=200011 | Both | This preprocessor symbol has the form *YYYYMM* where *YYYY* is the year and *MM* is the month of the OpenMP Fortran specification supported. This preprocessor symbol can be used in both fpp and the Fortran compiler conditional compilation. It is available only `-openmp` is specified. |
| `_PRO_INSTRUMENT` | Off | Both | Defined when `-prof_gen` is specified. |

# Defining Preprocessor Symbols

You can use the -D option to define the symbol names to be used during preprocessing. This option performs the same function as the #define preprocessor directive. The format of this option is:

*-Dname[=value]*

where:

- *name* is the name of the symbol to define
- *value* specifies an optional *value* to substitute for *name*

If you do not enter a *value*, *name* is set to 1. The *value* should be enclosed in quotation marks if it contains spaces or special characters.

Preprocessing replaces every occurrence of *name* with the specified *value*. For example, to define a symbol called *SIZE* with the *value* 100, use the following command:

**ifort -fpp -D*SIZE*=100 prog1.f**

Preprocessing replaces all occurrences of *SIZE* with the specified value (100) before passing the preprocessed source code to the compiler. Assume that the program contains this declaration:

```
REAL VECTOR(SIZE)
```

In the code sent to the compiler, the value 100 replaces *SIZE* in this declaration, and in all other occurrences of the name *SIZE*.

## Suppressing Preprocessor Symbols

You can use the `-U` option to suppress an automatic definition of a preprocessor symbol. This option suppresses any symbol definition currently in effect for the specified name. The `-U` option performs the same function as an `#undef` preprocessor directive.

# Redirecting Command-Line Output to Files

For programs that display a lot of text, consider redirecting text that is usually displayed on `stdout` to a file. Displaying a lot of text will slow down execution; scrolling text in a terminal window on a workstation can cause an I/O bottleneck (increased elapsed time) and use more CPU time.

The following commands show how to run the program more efficiently by redirecting output to a file and then displaying the program output:

```
myprog > results.lis
more results.lis
```

Redirecting output from the program will change the times reported because of reduced screen I/O.

# Creating, Running, and Debugging an Executable Program

The example below shows a sample Fortran main program using free source form that uses a module and an external subprogram.

The function CALC_AVERAGE is contained in a separate file and depends on the module ARRAY_CALCULATOR for its interface block.

The USE statement accesses the module ARRAY_CALCULATOR. This module contains the function declaration for CALC_AVERAGE.

The 5-element array is passed to the function CALC_AVERAGE, which returns the value to the variable AVERAGE for printing.

The example is:

```
! File: main.f90
! This program calculates the average of five numbers
  PROGRAM MAIN
    USE ARRAY_CALCULATOR
    REAL, DIMENSION(5) :: A = 0
    REAL :: AVERAGE
    PRINT *, 'Type five numbers: '
    READ  (*,'(F10.3)') A
    AVERAGE = CALC_AVERAGE(A)
    PRINT *, 'Average of the five numbers is: ', AVERAGE
  END PROGRAM MAIN
```

The example below shows the module referenced by the main program. This example program shows more Fortran 95/90 features, including an interface block and an assumed-shape array:

```
! File: array_calc.f90.
! Module containing various calculations on arrays.
  MODULE ARRAY_CALCULATOR
    INTERFACE
      FUNCTION CALC_AVERAGE(D)
        REAL :: CALC_AVERAGE
        REAL, INTENT(IN) :: D(:)
      END FUNCTION CALC_AVERAGE
    END INTERFACE
  ! Other subprogram interfaces...
  END MODULE ARRAY_CALCULATOR
```

The example below shows the function declaration CALC_AVERAGE referenced by the main program:

```
! File: calc_aver.f90.
! External function returning average of array.
  FUNCTION CALC_AVERAGE(D)
    REAL :: CALC_AVERAGE
    REAL, INTENT(IN) :: D(:)
    CALC_AVERAGE = SUM(D) / UBOUND(D, DIM = 1)
  END FUNCTION CALC_AVERAGE
```

# Commands to Create a Sample Program

During the early stages of program development, the sample program files shown above might be compiled separately and then linked together, using the following commands:

```
ifort -c array_calc.f90
ifort -c calc_aver.f90
ifort -c main.f90
ifort -o calc main.o array_calc.o calc_aver.o
```

In this sequence of commands:

- The `-c` option prevents linking and retains the `.o` files.
- The first command creates the files `array_calculator.mod` and `array_calc.o` (the name in the MODULE statement determines the name of module file `array_calculator.mod)`. Module files are written into the current working directory.
- The second command creates the file `calc_aver.o`.
- The third command creates the file `main.o` and uses t0he module file `array_calculator.mod`.
- The last command links all object files into the executable program named `calc`. To link files, use the `ifort` command instead of the `ld` command.

The order in which the file names are specified is significant. This `ifort` command:

- Compiles the file `array_calc.f90`, which contains the module definition, and creates its object file and the file `array_calculator.mod`.
- Compiles the file `calc_aver.f90`, which contains the external function CALC_AVERAGE.
- Compiles the file `main.f90` (main program). The USE statement references the module file `array_calculator.mod`.
- Uses `ld` to link the main program and all object files into an executable program file named `calc`.

# Running the Sample Program

If your path definition includes the directory containing `calc`, you can run the program by simply entering its name:

```
calc
```

When running the sample program, the PRINT and READ statements in the main program result in the following dialogue between user and program:

```
Type five numbers:
55.5
```

```
4.5
3.9
9.0
5.6
```
Average of the five numbers is:    15.70000

## Debugging the Sample Program

To debug a program with the debugger, compile the source files with the `-g` option to request additional symbol table information for source line debugging in the object and executable program files. The following `ifort` command also uses the `-o` option to name the executable program file `calc_debug`:

**ifort -g -o calc_debug array_calc.f90 calc_aver.f90 main.f90**

See also Debugging Overview and related sections.

# Creating Shared Libraries

To create a shared library from a Fortran source file, process the files using the `ifort` command:

- You must specify the `-shared` option to create the `.so` file.
- You can specify the `-o output` option to name the output file.
- If you omit the `-c` option, you will create a shared library (`.so` file) directly from the command line in a single step.
  If you also omit the `-o output` option, the file name of the first Fortran file on the command line is used to create the file name of the `.so` file. You can specify additional options associated with shared library creation.
- If you specify the `-c` option, you will create an object file (`.o` file) that you can name with the `-o` option. To create a shared library, process the `.o` file with `ld`, specifying certain options associated with shared library creation.
- When building shared libraries on Itanium-based systems, you must specify the `-fpic` option for the compilation of each object file included in the shared library. If this option is not used, the linker will probably emit an error message like `@gprel relocation against dynamic symbol`.

## Creating a Shared Library with a Single ifort Command

You can create a shared library (`.so`) file with a single `ifort` command:

**ifort -shared  octagon.f90**

The `-shared` option is required to create a shared library. The name of the source file is `octagon.f90`. You can specify multiple source files and object files.

The `-o` option was omitted, so the name of the shared library file is `octagon.so`.

Since you omitted the `-c` option, you do not need to specify the standard list of Fortran libraries.

# Creating a Shared Library with ifort and ld Commands

You first must create the `.o` file, such as `octagon.o` in the following example:

```
ifort -c octagon.f90
```

The file `octagon.o` is then used as input to the `ld` command to create the shared library named `octagon.so`:

```
ld -shared -no_archive octagon.o \
    -lifport -lifcoremt -limf -lm -lirc -lcxa \
    -lunwind -lpthread -lc
```

Note the following:

- The `-shared` option is required to create a shared library.
- The `-no_archive` option indicates that `ld` should not search archive libraries to resolve external names (only shared libraries).
- The name of the object file is `octagon.o`. You can specify multiple object (`.o`) files.
- The `-libport` and subsequent options are the standard list of libraries that the `ifort` command would have otherwise passed to `ld`. When you create a shared library, all symbols must be resolved.

It is probably a good idea to look at the output of the `-dryrun` command to find the names of all the libraries used so you can specify them correctly.

You can use the `-Qoption` command to pass options to `ld`.

See also the `ld(1)` reference page.

# Shared Library Restrictions

When creating a shared library with `ld`, be aware of the following restrictions:

- Shared libraries must not be linked with archive libraries.
  When creating a shared library, you can only depend on other shared libraries for resolving external references. If you need to reference a routine that currently resides in an archive library, either put that routine in a separate shared library or include it in the shared library being created. You can specify multiple object (`.o`) files when creating a shared library.
  To put a routine in a separate shared library, obtain the source or object file for that routine, recompile if necessary, and create a separate shared library. You can specify an object file when recompiling with the `ifort` command or when creating the shared library with the `ld` command.
  To include a routine in the shared library being created, put the routine (source or object file) with other source files that make up the shared library and recompile if necessary.
  Then create the shared library, making sure that you specify the file containing that routine either during recompilation or when creating the shared library. You can specify an object file when recompiling with the `ifort` command or when creating the shared library with the `ld` command.
- When creating shared libraries, all symbols must be defined (resolved). Because all symbols must be defined to `ld` when you create a shared library, you must specify the shared libraries on the `ld` command line, including all standard Intel Fortran libraries (unless you use the `-Qoption` command). The list of standard Intel Fortran libraries might be specified by using the `-lstring` option.

## Installing Shared Libraries

Once the shared library is created, it must be installed for private or system-wide use before you run a program that refers to it:

- To install a private shared library (when you are testing, for example), set the environment variable `LD_LIBRARY_PATH`, as described in `ld(1)`.
- To install a system-wide shared library, place the shared library file in one of the standard directory paths used by `ld`. See `ld(1)`.

# Allocating Common Blocks

You can use the `-dyncom` (dynamic common) option to control the allocation of common blocks at run time.

This option designates a common block to be dynamic. The space for its data is allocated at run time rather than compile time. On entry to each routine containing a declaration of the dynamic common block, a check is made of whether space for the common block has been allocated. If the dynamic common block is not yet allocated, space is allocated at the check time.

The following command-line example specifies the dynamic common option with
the names of the common blocks to be allocated dynamically at run time:

```
ifort -dyncom "blk1,blk2,blk3" test.f
```

where `blk1`, `blk2`, and `blk3` are the names of the common blocks to be made
dynamic.

## Guidelines for Using the -dyncom Option

The following are some limitations that you should be aware of when using the `-dyncom` option:

- An entity in a dynamic common cannot be initialized in a DATA statement.
- Only named common blocks can be designated as dynamic COMMON.
- An entity in a dynamic common block must not be used in an EQUIVALENCE expression with an entity in a static common block or a DATA-initialized variable.

## Why Use a Dynamic Common Block?

A main reason for using dynamic common blocks is to enable you to control the
common block allocation by supplying your own allocation routine. To use your
own allocation routine, you should link it ahead of the Fortran run-time library.
This routine must be written in the C language to generate the correct routine
name.

The routine prototype is:

```
void _FTN_ALLOC(void **mem, int *size, char *name);
```

*where*

- *`mem` is the location of the base pointer of the common block which must be set by the routine to point to the block of memory allocated.*
- *`size` is the integer number of bytes of memory that the compiler has determined are necessary to allocate for the common block as it was declared in the program. You can ignore this value and use whatever value is necessary for your purpose.*

### 📒 *Note*

*You must return the size in bytes of the space you allocate. The library
routine that calls `_FTN_ALLOC()` ensures that all other occurrences of this*

31

*common block fit in the space you allocated. Return the size in bytes of the space you allocate by modifying size.*

- `name` *is the name of the common block being dynamically allocated.*

## Allocating Memory to Dynamic Common Blocks

*The run-time library routine, f90_dyncom, performs memory allocation. The compiler calls this routine at the beginning of each routine in a program that contains a dynamic common block. In turn, this library routine calls _FTN_ALLOC() to allocate memory. By default, the compiler passes the size in bytes of the common block as declared in each routine to f90_dyncom, and then on to _FTN_ALLOC(). If you use the nonstandard extension having the common block of the same name declared with different sizes in different routines, you might get a run-time error depending on the order in which the routines containing the common block declarations are invoked.*

*The Fortran run-time library contains a default version of _FTN_ALLOC(), which simply allocates the requested number of bytes and returns.*

## Compiler Options

# Compiler Options Overview

*See these topics:*

*Compiler Options Details*

*Compiler Directives Related to Options*

*Code Generation Options*

*Compatibility Options*

*Compilation Diagnostics Options*

*Data Options*

*External Procedures Options*

*Floating-Point Options*

*Language Options*

# *Compiler Options Details*

*Options to the `ifort` command affect how the compiler processes a file in conjunction with the file name suffix. The simplest form of the `ifort` command is often sufficient.*

## **Styles of Compiler Options**

*Some options consist of two words separated by a space, while others may have words joined by an underscore ( _ ). Most options can be abbreviated, usually to four characters or more. For example, you can abbreviate `-check output_conversion` to `-check out.`*

*Intel Fortran has four styles of compiler options:*

- *The phrase `no` precedes the option to disable it. This style was used in Compaq\* Fortran. Example: `logo` and `nologo`*
- *A hyphen following the option disables it. This style was used in previous Intel Fortran compilers. Example: `-prefetch` and `-prefetch-`*
- *A `no` or `no-` in the middle of the option disables it. Example: `-falias` and `-fno-alias`*
- *The option has an `n` parameter (specifying a number) and is disabled by setting `n` equal to 0.*

**Note**

*If there are enabling and disabling versions of options on the command line, or two versions of the same option, the last one takes precedence.*

## **Using Multiple ifort Commands**

*If you compile parts of your program by using multiple `ifort` commands, options that affect the execution of the program should be used consistently for all compilations, especially if data is shared or passed between procedures. For example:*

- *The same data alignment needs to be used for data passed or shared by module definition (such as user-defined structures) or common block. Use the same version of the `-align` option for all compilations.*
- *The program might contain INTEGER, LOGICAL, REAL, COMPLEX, or DOUBLE PRECISION declarations without a kind parameter or size specifier that is passed or shared by module definition or common block. You must consistently use the options that control the size of such numeric data declarations.*

## Using the OPTIONS Statement to Override Options

*You can override some options specified on the command line by using the OPTIONS statement in your Fortran source program. The options specified by the OPTIONS statement affect only the program unit where the statement occurs.*

## Getting Help on Options

*For help, enter `-help` on the command line, which displays brief information about all the command-line options.*

# Compiler Directives Related to Options

*Some compiler directives and compiler options have the same effect, as shown in the table below. However, compiler directives can be turned on and off throughout a program, while compiler options remain in effect for the whole compilation unless overridden by a compiler directive.*

*Compiler directives and equivalent command-line compiler options are:*

| Compiler Directive | Equivalent Command-Line Compiler Option |
|---|---|
| DECLARE | -warn declarations |
| NODECLARE | -warn nodeclarations |
| DEFINE *symbol* | -D*name* |

| FIXEDFORMLINESIZE:*option* | `-extend_source [`*option*`]` |
|---|---|
| FREEFORM | `-free` **or** `-nofixed` |
| NOFREEFORM | `-nofree` **or** `-fixed` |
| INTEGER:*option* | `-integer_size` *option* |
| PACK:*option* | `-align [`*option*`]` |
| REAL:*option* | `-real_size` *option* |
| STRICT | `-warn stderrors` **with** `-stand` |
| NOSTRICT | `-warn nostderrors` |

*Note that the compiler directive names above are specified using the prefix `!DEC$` followed by a space.  For example: `!DEC$ NOSTRICT`*

### 🗒 *Note*

*The prefix `!DEC$` is normally used. `!DEC$` works for both fixed-form and free-form source. You can also use these alternative prefixes for fixed-form source only: `cDEC$,CDEC$,*DEC$,cDIR$,CDIR$, *DIR$,` and `!MS$`.*

# Code Generation Options

*The code generation options let you specify how code should be generated.*

## Descriptions of Code Generation Options

### -[no]recursive

*Default: `-norecursive`*

*Compiles all procedures (functions and subroutines) for possible recursive execution. When `-recursive` is specified, the `-auto` option is also set.*

### -[no]reentrancy [keyword]

*Default: `-noreentrancy`*

*Generates reentrant code that supports a multithreaded application. The `keyword` can be:*

- *`none`   Same as `-noreentrancy`.  Tells the Intel Fortran run-time library (RTL) that the program will not be relying on threaded or asynchronous reentrancy. Therefore, the RTL will not guard against such interrupts inside its own critical regions.*

- *async*
  Tells the RTL that the program may contain asynchronous handlers that could call the RTL. This causes the RTL to guard against asynchronous interrupts inside its own critical regions.
- *threaded*
  Tells the RTL that the program is multithreaded. This causes the RTL to use thread locking to guard its own critical regions.
  Specifying *-threads* sets *-reentrancy threaded,* since multithreaded code must be reentrant.
  Specifying *-reentrancy* is equivalent to specifying *-reentrancy threaded.*

### -sox[-] (IA-32 systems only)

Default: *-sox-*

Enables saving of the compiler options and version in the executable.

This option has no effect in Itanium®-based systems.

# Compatibility Options

The compatibility options let you specify how to make your source files and data files compatible with older Fortran versions or other operating systems, such as big endian unformatted data files, OpenVMS* systems run-time behavior, and Microsoft* Fortran PowerStation.

**See Also**

*Data Options*

*Language Options*

## Descriptions of Compatibility Options

### -1

Default: Off

Alternate syntax: *-onetrip*

Specifies that the compiler should execute DO loops at least once. See also *-[no]f66.*

### -assume [no]bscc

*Default:* `-assume nobscc`

*Alternate syntax:* `-nbs` *is the same as* `-assume bscc`

*Tells the compiler to treat the backslash character (\) as a C-style control (escape) character in character literals. The default,* `-assume nobscc` *("assume no BackSlashControlCharacters"), tells the compiler to treat the backslash character as a normal character instead of a control character in character literals.*

*This option is useful when transferring programs from non-UNIX\* environments, such as OpenVMS\*.*

### -convert

*Default: None.*

*Specifies the format of unformatted files containing numeric data. Possible values are:*

- *-convert big_endian*
- *-convert cray*
- *-convert ibm*
- *-convert little_endian*
- *-convert native*
- *-convert vaxg*
- *-convert vaxd*

*See* Supported Native and Nonnative Numeric Formats*.*

### -[no]f77rtl

*Default:* `-nof77rtl`

*Specifies the use of FORTRAN 77 run-time behavior. If you use the default value (*`-nof77rtl`*), Intel Fortran run-time behavior is used.*

*Specifying this option controls control the following run-time behavior:*

- *When the unit is not connected to a file, some INQUIRE specifiers will return different values:*
  *NUMBER returns 0*
  *ACCESS returns 'UNKNOWN'*

> *BLANK returns 'UNKNOWN'*
> *FORM returns 'UNKNOWN'*

- *List-directed input for character strings must be delimited by apostrophes or quotation marks, or an error will result.*
- *When processing NAMELIST input:*
  *Column 1 of each record is skipped*
  *The '$' or '&' that appears prior to the group-name must appear in column 2 of the input record*

## -fpscomp all and -fpscomp none

*Default: `-fpscomp libs`*

*Specifies that all the `-fpscomp` options for compatibility with Microsoft\* Fortran PowerStation should be used. The default value specifies that the PowerStation portability library should be passed to the linker.*

*`-fpscomp none` specifies that no options for Fortran PowerStation compatibility should be used.*

## -fpscomp [no]filesfromcmd

*Default: `-fpscomp nofilesfromcmd`*

*Specifies Microsoft\* Fortran PowerStation behavior when the OPEN statement file specifier is blank. This option looks in the command-line arguments for unspecified filenames on an OPEN(. . . FILE=' ', . . .) and prompts for filenames at the terminal console.*

## -fpscomp [no]general

*Default: `-fpscomp nogeneral`*

*Specifies that Microsoft\* Fortran PowerStation semantics should be used where differences exist between Intel Fortran and PowerStation.*

## -fpscomp [no]ioformat

*Default: `-fpscomp noioformat`*

*Specifies Microsoft\* Fortran PowerStation semantic conventions and record formats for list-directed formatted I/O and unformatted I/O.*

## -fpscomp [no]libs

*Default: `-fpscomp libs`*

*Specifies that the PowerStation portability library should be passed to the linker.*

## -fpscomp [no]logicals

*Default: `-fpscomp nologicals`*

*Specifies that Microsoft\* Fortran PowerStation representation of LOGICAL values will be used.*

## -prof_format_32

*Default: Off*

*Produces profile data with 32-bit counters. The default is to produce profile data with 64-bit counters to handle large numbers of events.*

*This option allows compatibility with earlier compilers.*

## -vms

*Default: Off*

*Causes the run-time system to behave like HP Fortran on OpenVMS Alpha systems and VAX\* systems (VAX FORTRAN\*) in the following ways:*

- *Certain defaults*
  *In the absence of other options, -`vms` sets the defaults as `-check format` and `-check output_conversion`.*
- *Alignment*
  *The `-vms` option does not affect the alignment of fields in records or items in common blocks. Use `-align norecords` to pack fields of records on the next byte boundary for compatibility with HP Fortran on OpenVMS systems.*
- *Carriage control default*
  *If `-vms -ccdefault default` is specified, carriage control defaults to FORTRAN if the file is formatted and the unit is connected to a terminal.*
- *INCLUDE qualifiers*
  */LIST and /NOLIST are recognized at the end of the file name in an INCLUDE statement at compile time.*
  *If the file name in the INCLUDE statement does not specify the complete path, the path used is the current directory.*
  *Note that if `-vms` is not specified, the path used is the directory where the file that contains the INCLUDE statement resides.*

- *Quotation mark character*
  *A quotation mark (") character is recognized as starting an octal constant ("0..7) instead of a character literal ("...").*
- *Deleted records in relative files*
  *When a record in a relative file is deleted, the first byte of that record is set to a known character (currently ' @ ' ). Attempts to read that record later result in ATTACCNON errors. The rest of the record (the whole record, if –vms is not specified) is set to nulls for unformatted files and spaces for formatted files.*
- *ENDFILE records*
  *When an ENDFILE is performed on a sequential unit, an actual 1-byte record containing a Ctrl/Z is written to the file. If –vms is not specified, an internal ENDFILE flag is set and the file is truncated.*
  *The –vms option does not affect ENDFILE on relative files: these files are truncated.*
- *Implied logical unit numbers*
  *The –vms option enables Intel Fortran to recognize certain environment variables at run time for ACCEPT, PRINT, and TYPE statements and for READ and WRITE statements that do not specify a unit number (such as READ (*,1000)).*
- *Treatment of blanks in input*
  *The –vms option causes the defaults for the keyword BLANK in OPEN statements to become ' NULL ' for an explicit OPEN and ' ZERO ' for an implicit OPEN of an external or internal file. For more information, see the description of the OPEN statement.*
- *OPEN statement effects*
  *Carriage control defaults to FORTRAN if the file is formatted, and the unit is connected to a terminal (checked by means of isatty(3) ). Otherwise, carriage control defaults to LIST.*
  *The –vms option affects the record length for direct access and relative organization files. The buffer size is increased by 1 to accommodate the deleted record character.*
- *Reading deleted records and ENDFILE records*
  *The run-time direct access READ routine checks the first byte of the retrieved record. If this byte is ' @ ' or NULL ("\0"), then an ATTACCNON error is returned.*
  *The run-time sequential access READ routine checks to see if the record it just read is one byte long and contains a Ctrl/Z. If this is true, it returns EOF.*

# *Compilation Diagnostics Options*

*The compilation diagnostics options let you specify the kinds of diagnostic messages (warnings and errors) you want to receive.*

# *Descriptions of Compilation Diagnostics Options*

## *-e90 or -e95*

*Default: Off*

*Alternate syntax: `-w90` or `-w95`*

*Issues errors for nonstandard Fortran 90 (`-e90`) or nonstandard Fortran 95 (`-e95`). This option issues compile-time errors for language elements that are not standard in the Fortran language that can be identified at compile time.00*

*See also `-[no]stand`.*

## *-[no]error_limit n*

*Default: `-error_limit 30`*

*Specifies the maximum number of error-level or fatal-level compiler errors allowed for a given file before compilation aborts. If you specify `-noerror_limit` on the command line, there is no limit on the number of errors that are allowed. If the maximum number of errors is reached, a warning message is issued and the next file (if any) on the command line is compiled.*

## *-openmp_report{0|1|2}*

*Default: Off. `-openmp_report1` is the default if `-openmp_report` is specified without an argument.*

*Specifies the OpenMP parallelizer's diagnostic level, where `n` is:*

- *`0` No information*
- *`1` Loops, regions, and sections parallelized*
- *`2` Same as `1` plus master construct, single construct, and so forth*

*For more information, see "Parallelization with OpenMP* Overview" (and related sections) in the User's Guide Volume II: Optimizing Applications.*

## *-par_report{0|1|2|3}*

*Default: Off. `-par_report1` is the default if `-par_report` is specified without an argument.*

*Specifies the autoparallelizer's diagnostic level, where `n` is:*

- *0 No information*
- *1 Loops successfully parallelized*
- *2 Loops successfully and unsuccessfully parallelized*
- *3 Same as 2 plus dependency information*

*See also these topics in Volume II:*

*Auto-Parallelization Overview*
*Auto-Parallelization: Enabling, Options, Directives, and Environment Variables*

## -std, -std90, -std95

*Default: Off ((no messages are issued)*

*Alternate syntax: `-[no]stand` or `-w90` or `-stand90` (for Fortran 90) or `-w95` or `-stand95` (for Fortran 95)*

*`-std` and `-stand` and `-std95` and `-stand95` (which are equivalent) warn for nonstandard Fortran 95. `-std90` and `-stand90` (which are equivalent) warn for nonstandard Fortran 90.*

*This option issues compile-time messages for language elements that are not standard in the Fortran language that can be identified at compile time.*

*`-w90` and `-w95` turn off warnings for nonstardard Fortran for Fortran 90 and Fortran 95, respectively.*

*`-stand` is set if you specify `-warn stderrors`.*

## -vec_report{0|1|2|3|4|5} (IA-32 systems only)

*Default: Off. `-vec_report1` is the default if `-vec_report` is specified without an argument.*

*Specifies the vectorizer's diagnostic level, where `n` is:*

- *0 No information*
- *1 Indicate vectorizer loops*
- *2 Indicate vectorizer and nonvectorizer loops*
- *3 Indicate vectorizer loops plus dependence information*
- *4 Indicate nonvectorized loops*
- *5 Indicate nonvectorized loops plus the reason why they were not vectorized*

*For more information, see "Vectorization Overview" (and related sections) in the User's Guide Volume II: Optimizing Applications.*

### -warn all or -warn none or -nowarn

*Default: Custom (individually specified).*

*Specifies the compiler diagnostics level. Choices are:*

- *`-warn all` (show all diagnostics)*
- *`-warn none` (show no diagnostics)*

*Specifying `-warn all` requests all possible warning messages, but does not set `-warn errors` or `-warn stderrors`. To enable all the additional checking to be performed and force the severity of the diagnostics to be severe enough to not generate an object file, specify `-warn all -warn errors` or `-warn all -warn stderrors`.*

*Specifying `-warn` is the same as specifying `-warn all`.*

*Specifying `-nowarn` is the same as specifying `-warn none`.*

### -warn [no]alignments

*Default: `-warn alignments`*

*Issues warning messages for data that is not naturally aligned.*

### -warn [no]declarations

*Default: `-warn nodeclarations`*

*Issues an error message for any undeclared symbols. This option makes the default type of a variable undefined (IMPLICIT NONE) rather than using the implicit Fortran rules. See also `-u`.*

### -warn [no]errors

*Default: `-warn noerrors`*

*Treats all warnings as errors by changing the severity of all warning diagnostics into error diagnostics, including standards warnings.*

### -warn [no]general

*Default: `-warn general`*

*Alternate syntax: `-W1` (to display all warnings) or `-W0` or `-w` (to suppress all warnings)*

*Displays all informational-level and warning-level diagnostic messages from the compiler.*

*Use `-warn nogeneral` or `-nowarn` or `-W0` or `-w` to suppress all warnings.*

## -warn [no]ignore_loc

*Default: `-warn noignore_loc`*

*Issues warning messages when %LOC is stripped from an argument.*

## -warn [no]stderrors

*Default: `-warn nostderrors`*

*Treats warnings about Fortran standards violations as errors, not warnings.*

*Specifying `-warn stderrors` sets `-stand f95.`*

*If you want to make Fortran 90 standards violations become errors, set this option as well as `-stand f90.`*

## -warn [no]truncated_source

*Default: `-warn notruncated_source`*

*Issues warning messages when reading a source line with a statement field that exceeds the maximum column width in fixed-format source files. The maximum column width for fixed-format files is 72, 80, or 132, depending on the setting of the `-extend_source` option. The `-warn truncated_source` option has no effect on truncation; lines that exceed the maximum column width are always truncated. The `-warn truncated_source` option does not apply to free-format source files.*

## -warn [no]uncalled

*Default: `-warn uncalled`*

*Issues warning messages when a statement function is never called.*

### *-warn [no]unused*

*Default: `-warn nounused`*

*Issues warning messages for variables that are declared but never used.*

### *-warn [no]usage*

*Default: `-warn usage`*

*Alternate syntax: `-cm` (which is equivalent to `-warn nousage`)*

*Suppresses messages about questionable programming practices.*

*Questionable programming practices, although allowed, often are the result of programming errors. For example, the default value, `-warn usage`, detects a continued character or Hollerith literal whose first part ends before the statement field ends and appears to end with trailing spaces.*

# Data Options

The data options let you specify rules for how your Fortran data is treated by the compiler, optimizer, and code generator.

**See Also**
Compatibility Options
Language Options

## Descriptions of Data Options

### -[no]align

Default: `-align`

Analyzes and reorders memory layout for variables and arrays. See also `-align recnbyte`.

### -align none

Default: Off

Tells the compiler not to add padding bytes anywhere in common blocks or structures.

## -align [no]commons or -align [no]dcommons

Default: `-align nocommons` or `-align nodcommons`

Aligns the data items of all common blocks on natural boundaries up to 4 bytes (`-align commons`) or 8 bytes (`-align dcommons`) instead of the default byte boundary by adding padding bytes.

If your command line includes the `-stand` option, then the compiler ignores `-align dcommons`.

## -align recnbyte

Default: `-align rec8byte`, `-Zp8`. `-align rec8byte` is the same as specifying `-align records`.

Alternate syntax: `-Zp{1|2|4|8|16}`

Specifies alignment constraint for structures on 1-, 2-, 4-, 8-, or 16-byte boundaries.

Aligns fields of records and components of derived types on the smaller of the size boundary specified ($n$ can be 1, 2, 4, 8, or 16) or the boundary that will naturally align them.

Specifying `-align rec`$n$`byte` does not affect whether common blocks are naturally aligned or packed.

| This option: | Is the same as this option: |
|---|---|
| `-Zp` | `-align records` or `-align rec8byte` |
| `-Zp1` | `-alignment norecords` or `-align rec1byte` |
| `-Zp2` | `-align rec2byte` |
| `-Zp4` | `-align rec4byte` |
| `-align` | `-Zp8` with `-align dcommons`, `-align all`, or `-align dcommons` and `-align records` |
| `-noalign` | `-Zp1`, `-align none`, or `-align nocommons` and `-align nodcommons` and `-align norecords` |
| `-align rec1byte` | `-align norecords` |
| `-align rec8byte` | `-align records` |

### -align [no]records

Default: `-align records`

Requests that components of derived types and fields of records be aligned on natural boundaries up to 8 bytes (for derived types with the SEQUENCE statement) by adding padding. See `-align sequence`.

The `-align norecords` option requests that components and fields be aligned on arbitrary byte boundaries, instead of on natural boundaries up to 8 bytes, with no padding.

### -align [no]sequence

Default: `-align nosequence`

Tells the compiler that components of derived types with the SEQUENCE attribute will obey whatever alignment rules are currently in use. The default alignment rules align unsequenced components on natural boundaries.

The default value of `-align nosequence` means that components of derived types with the SEQUENCE attribute will be packed, regardless of whatever alignment rules are currently in use.

If your command line includes the `-stand` option, then the compiler ignores `-align sequence`.

### -assume [no]byterecl

Default: `-assume nobyterecl`

Specifies the use of byte units for unformatted files. This option:

- Specifies that the units for an explicit OPEN statement RECL specifier value are in bytes.
- Forces the record length value returned by an INQUIRE by output list to be in byte units.

The default value, `-assume nobyterecl`, specifies that the units for RECL values with unformatted files are in four-byte (longword) units.

### -assume [no]dummy_aliases

Default: `-assume nodummy_aliases`

Alternate syntax: `-common_args`

Requires that the compiler assume that dummy (formal) arguments to procedures share memory locations with other dummy arguments or with variables shared through use association, host association, or common block use.
You only need to compile the called subprogram with `-assume dummy_aliases`.
The program semantics involved with dummy aliasing do not strictly obey the Fortran standard and they slow performance. Therefore, using the default value, `-assume nodummy_aliases`, will give the compiler better run-time performance. However, if a program depends on dummy aliasing and you do not specify `-assume dummy_aliases,` the run-time behavior of the program will be unpredictable. In such programs, the results will depend on the exact optimizations that are performed. In some cases, normal results will occur, but in other cases, results will differ because the values used in computations involving the offending aliases will differ.

For more information, see the information about the dummy aliasing assumption in the *User's Guide Volume II: Optimizing Applications.*

## -assume [no]protect_constants

Default: `-assume protect_constants`

Specifies that constants are read-only.

The `-assume noprotect_constants` option tells the compiler to pass a copy of a constant actual argument. As a result, the called routine can modify this copy, even though the Fortran standard prohibits such modification. The calling routine does not modify the constant.

The default, `-assume protect_constants`, results in passing of a constant actual argument. Any attempt to modify it may result in an error.

## -auto_scalar, -auto, and -save

Default: `-auto_scalar` (unless `-recursive` or `-openmp` is specified, in which case the default is `-auto`)

Alternate syntax for `-auto`: `-automatic` and `-nosave`

Alternate syntax for `-save`: `-noauto` and `-noautomatic`

Specifies where local variables are stored, if default local storage is not desired.

`-auto_scalar` causes allocation of scalar variables of intrinsic types INTEGER, REAL, COMPLEX, and LOGICAL to the stack.

`-auto_scalar` does not affect variables that appear in EQUIVALENCE or SAVE statements, or those that are in common blocks. `-auto_scalar` may provide a performance gain for your program, but if your program depends on variables having the same value as the last time the routine was invoked, your program may not function properly. Variables that need to retain their values across subroutine calls should appear in a SAVE statement.

`-auto` makes all local variables AUTOMATIC, causing all variables to be allocated on the stack, rather than in local static storage. It does not affect variables that have the SAVE attribute or appear in an EQUIVALENCE statement or in a common block.

`-save` saves all variables in static allocation except local variables within a recursive routine.

`-auto` might provide a performance gain for your program, but if your program depends on variables having the same value as the last time the routine was invoked, your program might not function properly.

## -double_size {64|128}

Default: `-double_size 64`

Defines DOUBLE PRECISION and DOUBLE COMPLEX declarations, constants, functions, and intrinsics.

`-double_size 64` defines DOUBLE PRECISION declarations, constants, functions, and intrinsics as REAL*8 and defines DOUBLE COMPLEX declarations, functions, and intrinsics as COMPLEX*16.

`-double_size 128` defines DOUBLE PRECISION declarations, constants, functions, and intrinsics as REAL*16 and (for Fortran 90 and 95) defines DOUBLE COMPLEX declarations, functions, and intrinsics as COMPLEX*32.

## -dyncom "blk1,blk2,..."

Enables dynamic allocation of the specified common blocks at run time.

Example: For common  blocks A, B, and C, use this syntax:

`-dyncom "a,b,c"`

## -integer_size {16|32|64}

Default: `-integer_size 32`

Alternate syntax: `-i{2|4|8},` where 2, 4, and 8 stand for the KIND of integer and logical variables
Specifies the default size (in bits) of integer and logical declarations, constants, functions, and intrinsics, where $n$ is 16, 32, or 64:

- $n$ is 16: Makes the default integer and logical variables 2 bytes long. INTEGER and LOGICAL declarations are treated as (KIND=2). Alternate syntax: `-i2`. See also INTEGER(KIND=2) Representation.
- $n$ is 32: Makes the default integer and logical variables 4 bytes long. INTEGER and LOGICAL declarations are treated as (KIND=4). Alternate syntax: `-i4`. See also INTEGER(KIND=4) Representation.
- $n$ is 64: Makes the default integer and logical variables 8 bytes long. INTEGER and LOGICAL declarations are treated as (KIND=8). Alternate syntax: `-i8`. See also INTEGER(KIND=8) Representation.

## -pg

Default: Off.

Compiles and links for function profiling with `gprof(1)`.

This is the same as specifying `-p` or `-qp`.

## -real_size {32|64|128}

Default: `-real_size 32`

Alternate syntax: `-r{8|16},` where 8 and 16 stand for the KIND of real variables.

Specifies the default size (in bits) of real and complex declarations, constants, functions, and intrinsics, where $n$ is 32, 64, or 128:

- $n$ is 32: Defines REAL declarations, constants, functions, and intrinsics as REAL(KIND=4) (SINGLE PRECISION) and defines COMPLEX declarations, constants, functions, and intrinsics as COMPLEX(KIND=4) (COMPLEX).
  Alternate syntax: None.
  **See Also**
  REAL(KIND=4) (REAL) Representation
  COMPLEX(KIND=4) (COMPLEX) Representation

- $n$ is 64: Defines REAL declarations, constants, functions, and intrinsics as REAL(KIND=8) (DOUBLE PRECISION) and defines COMPLEX declarations, constants, functions, and intrinsics as COMPLEX(KIND=8) (DOUBLE COMPLEX).
  Specifying `-real_size 64` causes intrinsic functions to produce a REAL(KIND=8) or COMPLEX(KIND=8) result instead of a REAL(KIND=4) or COMPLEX(KIND=4) result, unless the argument is explicitly typed as REAL(KIND=4) or COMPLEX(KIND=4), including CMPLX, FLOAT, REAL, SNGL, and AIMAG. For instance, references to the CMPLX intrinsic produce DCMPLX results (COMPLEX(KIND=8)), unless the argument to CMPLX is explicitly typed as REAL(KIND=4), REAL*4, COMPLEX(KIND=4), or COMPLEX*8. In this case the resulting data type is COMPLEX(KIND=4).
  Alternate syntax: `-r8` or `-autodouble`
  **See Also**
  REAL(KIND=8) (DOUBLE PRECISION) Representation
  COMPLEX(KIND=8) (DOUBLE COMPLEX) Representation
- $n$ is 128: Defines REAL declarations, constants, functions, and intrinsics as REAL(KIND=16) and defines COMPLEX declarations, constants, functions, and intrinsics as COMPLEX(KIND=16).
  Alternate syntax: `-r16`
  **See Also**
  REAL(KIND=16) Representation
  COMPLEX(KIND=16) Representation

## -safe_cray_ptr

Default: Off (assume that Cray pointers do alias other variables)

Requires that the compiler assume that Cray pointers do not alias (that is, do not specify sharing memory between) other variables.

Consider the following example:

```
pointer (pb, b)
pb = getstorage()
do i = 1, n
  b(i) = a(i) + 1
enddo
```

By default, the compiler assumes that `b` and `a` are aliased. To prevent such an assumption, specify the `-safe_cray_ptr` option, and the compiler will treat `b(i)` and `a(i)` as independent of each other.

However, if the variables are intended to be aliased with Cray pointers, using the `-safe_cray_ptr` option produces incorrect results. For the code example below, `-safe_cray_ptr` should not be used:

```
pointer (pb, b)
pb = loc(a(2))
   do i=1, n
      b(i) = a(i) +1
   enddo
```

### -zero[-]

Default: Off (`-zero-`)

Initializes to zero all local scalar variables of intrinsic type INTEGER, REAL, COMPLEX, or LOGICAL that are saved and not already initialized.

For this option to be effective, you must use `-save` on the command line or have variables in your code specifically marked as SAVE.

# External Procedures Options

The external procedures options let you specify how external procedures are called.

## Descriptions of External Procedures Options

### -assume [no]underscore

Default: `-assume nounderscore`

Alternate syntax: `-nus[,filename]`

Appends an underscore character to external user-defined names: the main program name, named COMMON, BLOCK DATA, global data names in MODULEs, and names implicitly or explicitly declared EXTERNAL. The name of blank COMMON remains _BLNK__, and Fortran intrinsic names are not affected.

If you want to specify a particular filename in which you don't want to append an underscore, use `-nus,filename`.

### -[no]mixed_str_len_arg

Default: `-nomixed_str_len_arg`

52

Tells the compiler that the hidden length passed for a character argument is to be placed immediately after its corresponding character argument in the argument list.

The default value places the hidden lengths in sequential order at the end of the argument list. When porting mixed-language programs that pass character arguments, either this option must be specified correctly or the order of hidden length arguments changed in the source code.

See also Programming with Mixed Languages Overview and related sections.

### -names keyword

Default: `-names lowercase`

Controls how the case sensitivity of letters in source code identifiers and external names is handled. This naming convention applies whether names are being defined or referenced. This option is useful in mixed-language programming.

`keyword` is one of the following:

- `uppercase:` Causes the compiler to ignore case differences in identifiers and to convert external names touppercase. Alternate syntax: `-uppercase`
- `lowercase:` Causes the compiler to ignore case differences in identifiers and to convert external names to lowercase. Alternate syntax: `-lowercase`
- `as_is:` Causes the compiler to distinguish case differences in identifiers and to preserve the case of external names.

Instead of using this option, consider using the ALIAS directive for the specific name needed.

# Floating-Point Options

The floating-point options let you specify how you want floating-point data to be treated.

See also Optimization Options.

## Descriptions of Floating-Point Options

### -assume [no]minus0

Default: `-assume nominus0`

Tells the compiler to use Fortran 95 standard semantics for the treatment of the IEEE* floating-point value -0.0 in the SIGN intrinsic, if the processor is capable of distinguishing the difference between -0.0 and +0.0, and to write a value of -0.0 with a negative sign on formatted output.

The default, `-assume nominus0`, tells the compiler to use Fortran 90/77 standard semantics in the SIGN intrinsic, to treat -0.0 and +0.0 as 0.0, and to write a value of -0.0 with no sign on formatted output.

## -[no]fltconsistency

Default: `-nofltconsistency`

Alternate syntax: `-mp`

Enables improved floating-point consistency during calculations.

This option limits floating-point optimizations and maintains declared precision. Floating-point operations are not reordered and the result of each floating-point operation is stored into the target variable rather than being kept in the floating-point processor for use in a subsequent calculation.

For example, the compiler can change floating-point division computations into multiplication by the reciprocal of the denominator. This change can alter the results of floating-point division computations slightly.

The default value, `-nofltconsistency`, provides better accuracy and run-time performance at the expense of less consistent floating-point results.

This option might slightly reduce execution speed.

See also "Improving/Restricting FP Arithmetic Precision" in *Volume II: Optimizing Applications*.

## -fp_port (IA-32 systems only)

Default: Off

Rounds floating-point results after floating-point operations, so rounding to user-declared precision happens at assignments and type conversions; this has some impact on speed.

The default is to keep results of floating-point operations in higher precision; this provides better performance but less consistent floating-point results.

See also "Floating-Point Arithmetic Precision for IA-32 Systems" in *Volume II: Optimizing Applications*.

## -[no]fpconstant

Default: `-nofpconstant`

Requests that a single-precision constant assigned to a double-precision variable be evaluated in double precision instead of single precision.

The Fortran standard requires that the constant be evaluated in single precision. Certain programs created for FORTRAN-77 compilers may show different floating-point results, because they rely on single-precision constants assigned to a double-precision variable to be evaluated in double precision.

In the following example, if you specify `/fpconstant`, identical values are assigned to `D1` and `D2`. If you omit `/fpconstant`, the compiler will obey the standard and assign a less precise value to D1:

```
REAL (KIND=8) D1, D2
DATA D1 /2.71828182846182/    ! REAL (KIND=4) value expanded
to double
DATA D2 /2.71828182846182D0/ ! Double value assigned to
double
```

## -fpen

Default: `-fpe3`

Specifies floating-point exception handling at run time for the main program. This includes whether exceptional floating-point values are allowed and how precisely run-time exceptions are reported. This option controls how the following exceptions are handled:

- When floating-point calculations result in a divide by zero, overflow, or invalid operation.
- When floating-point calculations result in an underflow.
- When a denormalized number or other exceptional number (positive infinity, negative infinity, or a NaN) is present in an arithmetic expression.

You can choose the following:

- `-fpe0` specifies: underflow gives 0.0; abort on other IEEE exceptions

- `-fpe3` specifies: produce NaN, signed infinities, and denormal results

On IA-32 systems, using `-fpe0` will slow run-time performance.

Many programs do not need to handle denormalized numbers or other exceptional values. On Itanium®-based systems, using `-fpe3` will slow run-time performance.

## -fpstkchk (IA-32 systems only)

Default: Off

Generates extra code after every function call to ensure that the floating-point (FP)stack is in the expected state.

By default, there is no checking. So when the FP stack overflows, a NaN value is put into FP calculations, and the program's results differ. Unfortunately, the overflow point can be far away from the point of the actual bug. The `-fpstkchk` option places code that would access-violate immediately after an incorrect call occurred, thus making it easier to locate these issues.

## -fr32 (Itanium®-based systems only)

Default: Off

Specifies that the use of high floating-point registers should be disabled.

## -ftz[-]

Default: Off (`-ftz-`) on IA-32 systems; off (`-ftz-`) on Itanium-based systems, except for optimization level `-O3`, in which case the default is on (`-ftz`)

Enables flush denormal results to zero. This option has effect only when compiling the main program.

## -IPF_flt_eval_method0 (Itanium®-based systems only)

Default: Off

Directs the compiler to evaluate the expressions involving floating-point operands in the precision indicated by the variable types declared in the program. By default, intermediate floating-point expressions are maintained in higher precision.

See also "Floating-Point Arithmetic Options for Itanium-Based Systems" in
*Volume II: Optimizing Applications.*

## -IPF_fltacc[-] (Itanium®-based systems only)

Default: Off

Disables optimizations that affect floating-point accuracy.

If the default setting is used (`-IPF_fltacc-`), the compiler might apply
optimizations that reduce floating-point accuracy.

You can use `-IPF_fltacc` or `-fltconsistency` to improve floating-point
accuracy, but at the cost of disabling some optimizations.

See also "Floating-Point Arithmetic Options for Itanium-Based Systems" in
*Volume II: Optimizing Applications.*

## -IPF_fma[-] (Itanium®-based systems only)

Default: Off

Enables the combining of floating-point multiply and add/subtract operations into
a single operation.

## -IPF_fp_speculationmode (Itanium®-based systems only)

Default: `-IPF_fp_speculationfast`

Enables floating-point operations in one of the following modes:

- `fast`    Speculate on floating-point operations
- `safe`    Speculate on floating-point operations only when it is safe
- `strict`  Disables speculation on floating-point operations.
- `off`     Same as strict.

See also "Floating-Point Arithmetic Options for Itanium-Based Systems" in
*Volume II: Optimizing Applications.*

## -mp1

Default: Off

Restricts floating-point precision to be closer to declared precision. This option
has some impact on speed, but less than the impact of `-mp`.

**-pc{32|64|80} (IA-32 systems only)**

Default: `-pc64`

Enables floating-point significand precision control. Possible values are:

- `-pc32`   Sets internal FPU precision to 24-bit significand
- `-pc64`   Sets internal FPU precision to 53-bit significand
- `-pc80`   Sets internal FPU precision to 64-bit significand

# Language Options

The language options let you specify semantics, syntax, and source file format.

**See Also**

Compatibility Options

Data Options

## Descriptions of Language Options

### -[no]altparam

Default: `-altparam` (alternate syntax is allowed)

Alternate syntax: `-[no]dps`

Allows alternate syntax (without parentheses) for PARAMETER statements.

The alternate syntax for PARAMETER statements is:

```
PARAMETER par1=exp1 [, par2=exp2] ...
```

This form does not have parentheses around the assignment of the constant to the parameter name. With this form, the type of the parameter is determined by the type of the expression being assigned to it and not by any implicit typing.

### -[no]d_lines

Default: `-nod_lines`

Alternate syntax: `-DD`

Specifies that lines in fixed-format files that contain a D in column 1 should be treated as source code, not comment lines.

## -[no]extend_source [size]

Default: `-noextend_source`, which implies 72 characters; if `-extend_source` is specified without a *size*, the default becomes `-extend_source:132`

Alternate syntax: `-{72|80|132}`

Specifies the column number used to end the statement field in fixed-form source files: 72, 80, or 132. When a size is specified, that will be the last column parsed as part of the statement field. Any columns after that will be treated as comments.

This option is valid only for fixed-form files, and it enables the `-fixed` option.

## -[no]f66

Default: `-nof66` (use current Fortran standards semantics)

Alternate syntax: `-66`
Specifies that the compiler should select FORTRAN-66 interpretations in cases of incompatibility. Differences include the following:

- DO loops are always executed at least once
- FORTRAN-66 EXTERNAL statement syntax and semantics are allowed
- If the OPEN statement STATUS specifier is omitted, the default changes to STATUS='NEW' instead of STATUS='UNKNOWN'
- If the OPEN statement BLANK specifier is omitted, the default changes to BLANK='ZERO' instead of BLANK='NULL'

## -[no]free or -[no]fixed

Default: File extension is used to determine format.

Alternate syntax: `-FR` is equivalent to `-free`; `-FI` is equivalent to `-fixed`.

Specifies the format of the Fortran source code. If this option is not specifed, the file extension determines the format:

- Files with an extension of `.f90`, `.F90`, or `.i90` are free-format source files.
- Files with an extension of `.f`, `.for`, `.FOR`, `.ftn`, or `.i` are fixed-format files.

**-openmp or -openmp_stubs**

Default: Off (disabled)

Specifies that OpenMP* directives should be processed. Options are:

- `-openmp`          Generate parallel code. If you use this option, multithreaded libraries are used, but fpp is not automatically invoked. When `-openmp` is specified, the `-auto` option is also set.
- `-openmp_stubs`   Generate sequential code. The OpenMP directives are ignored and a stub OpenMP library is linked.

**-[no]pad_source**

Default: `-nopad_source`

Specifies that fixed-form source lines shorter than the statement field width are to be padded with spaces to the end of the statement field. This affects the interpretation of character and Hollerith literals that are continued across source records.

The default value, `-nopad_source`, causes a warning message to be displayed if a character or Hollerith literal that ends before the statement field ends is continued onto the next source record. To suppress this warning message, specify the `-warn nousage` option.

Specifying `-pad_source` can prevent warning messages associated with `-warn usage`.

# Libraries Options

The libraries options let you specify libraries for your application.

## Descriptions of Libraries Options

### -no_cpprt

Default: `-cxxlib-icc` (use Intel C++ libraries)

Specifies that C++ run-time libraries should not be linked.

This option exists for GNU compatibility reasons, to disable the use of the cpp run-time libraries during link. There is no `-cpprt` or `-yes_cpprt` option.

**-nodefaultlibs**

Default: Off (include default libraries)

Specifies that standard libraries should be used when linking.

This option exists for GNU compatibility reasons. There is no `-defaultlibs`
option.

See also `-nostdlib`.

**-i_dynamic**

Default: Off

Instructs the linker to link Intel-provided libraries dynamically.

**-Ldir**

Default: Off

Instructs the linker to search *dir* for libraries.

**-[no]threads**

Default: `-nothreads`

Specifies whether or not multithreaded libraries should be linked against.

If you specify `-threads`, this sets the `-reentrancy threaded` option.

**-nostdlib**

Default: Off.

Specifies that standard libraries and startup files should be used when linking.

This option exists for GNU compatibility reasons. There is no `-stdlib` option.

**-shared**

Default: Off

Specifies that the compiler should build a dynamic shared object (DSO) instead
of an executable.

See also Creating Shared Libraries.

### -shared-libcxa

Links the Intel-provided `libcxa` C++ library dynamically.

By default, the `libcxa` library is linked dynamically. (All C++-related libraries supplied by Intel are linked in dynamically by default.) This option is useful when you are using the `-static` option and you want to override the effect of the `-static` option for the `libcxa` library.

This option has the opposite effect of `-static-libcxa`.

### -static-libcxa

Links the Intel-provided `libcxa` C++ library statically.

By default, the `libcxa` library is linked dynamically. (All C++-related libraries supplied by Intel are linked in dynamically by default.) Use this option to link `libcxa` statically, while still allowing the standard libraries to be linked in by the default behavior.

See also `-shared-libcxa` option.

### -static

Default: Off.

Alternative syntax: `-non_shared`

Prevents linking with shared libraries (`.so` files).

# Miscellaneous Options

These options are described in alphabetical order.

### -ansi_alias[-]

Default: `-ansi_alias`
Enables the compiler to assume that the program adheres to the Fortran 95 Standard type aliasability rules.

For example, an object of type real cannot be accessed as an integer. For complete information on the rules for data types and data type constants, see "Data Types, Constants, and Variables" in the *Language Reference*.

The option directs the compiler to assume the following:

- Arrays are not accessed out of arrays' bounds.
- Pointers are not cast to non-pointer types and vice-versa.
- References to objects of two different scalar types cannot alias. For example, an object of type integer cannot alias with an object of type real or an object of type real cannot alias with an object of type double precision.

If your program satisfies the above conditions, setting the `-ansi_alias` option will help the compiler optimize the program. However, if your program might not satisfy any of the above conditions, you must disable this option with `-ansi_alias-`, as it might cause the compiler to generate incorrect code.

## -assume cc_omp

Default: depends on whether `-openmp` is specified

Enables conditional compilation as defined by the OpenMP Fortran API. When `"!$space"` appears in free-form source or `"!$spaces"` appears in column 1 of fixed-form source, the rest of the line is accepted as a Fortran line.

If `-openmp` is specified, the default is `-assume cc_omp`; otherwise, the default is `-assume nocc_omp`.

## -assume none

Turns off all the `-assume` options.

## -nobss_init

Default: Off

Disables placement of zero-initialized variables in the BSS section.

By default, variables explicitly initialized with zeros are placed in the BSS section. By using this option, you can place any variables that are explicitly initialized with zeros in the DATA section if required.

There is no `-bss_init` option.

## -ccdefault keyword

Default: `-ccdefault default`

Specifies the type of carriage control used for units 6 and *.

The choices for `keyword` are:

- `default`   Specifies that the compiler is to use the default carriage-control setting. This choice can be affected by the `-vms` option:
  If `-vms -ccdefault default` is specified, carriage control defaults to `fortran` if the file is formatted, and the unit is connected to a terminal.
  If `-novms -ccdefault default` is specified, carriage control defaults to `list`.
- `fortran`   Specifies normal Fortran interpretation of the first character.
- `list`       Specifies one line feed between records.
- `none`       Specifies no carriage-control processing.

## -dryrun

Default: Off

Specifies that driver tool commands should be shown but not executed. See also `-v`.

## -dynamic-linkerfile

Default: Off

Specifies a dynamic linker (`file`) instead of the default.

## -fpic or -fPIC

Default: Off

Specifies that position-independent code should be generated.

Specifies full symbol preemption. Global symbol definitions as well as global symbol references get default (that is, preemptable) visibility unless explicitly specified otherwise.

See also Creating Shared Libraries.

## -fvisibility=keyword and -fvisibility-keyword=file

Specifies the default visibility for global symbols (`-fvisibility=`*`keyword`*) or
specifies the visibility for symbols that are in a file (`-fvisibility-`
*`keyword=file`*). (This second form overrides the first form).

The *`keyword`* specifies what the visibility is set to. Visibility can be set to any of
the following:

- `default` - Other components can reference the symbol, and the symbol
  definition can be overridden (preempted) by a definition of the same name
  in another component.
- `extern` - The symbol is treated as though it is defined in another
  component.  It also means that the symbol can be overridden by a
  definition of the same name in another component.
- `hidden` - Other components cannot directly reference the symbol.
  However, its address might be passed to other components indirectly.
- `internal` - The symbol cannot be referenced outside its defining
  component, either directly or indirectly.
- `protected` - Other components can reference the symbol, but it cannot
  be overridden by a definition of the same name in another component.

The *`file`* is the pathname of a file containing the list of symbols whose visibility
you want to set. The symbols are separated by whitespace (spaces, tabs, or
newlines).

## -g

Default: Off

Generates symbolic debugging information and line numbers in the object file for
use by debuggers.

## -help

Displays brief information about all the command-line options.

## -inline_debug_info

Default: Off

Keeps the source position of inline code instead of assigning the call-site source
position to inlined code.

## -[no]logo

Default: `-logo` (startup banner is displayed)

Displays the startup banner.

This option can be placed anywhere on the command line.

The startup banner displays the following information:

- `ID`: unique identification number for the compiler
- `x.y.z`: version of the compiler
- `years`: years for which the software is copyrighted

## -nofor_main

Default: Off

Specifies that the main program is not written in Fortran. For example, if the main program is written in C and calls an Intel Fortran subprogram, specify `-nofor_main` when compiling the program with the `ifort` command. Specifying `-nofor_main` prevents linking `for_main.o` into programs. This is a link-time switch.

If you omit `-nofor_main`, the main program must be a Fortran program.

## -noinclude

Default: Off

Prevents the compiler from searching in `/usr/include` for files specified in an INCLUDE statement.

You can specify the `-I`*dir* option along with this option. This option does not affect `cpp(1)` behavior, and is not related to the Fortran 95 and 90 USE statement.

## -[no]pad

Default: `-nopad`

Enables the changing of the variable and array memory layout.

The `-pad` option is effectively not different from `-align` when applied to structures and derived types. However, the scope of `-pad` is greater because it applies also to common blocks, derived types, sequence types, and structures.

## -prec_div (IA-32 systems only)

Default: Off

Enables improved precision of floating-point divides. Has a slight impact on speed.

### -rcd (IA-32 systems only)

Default: Off

Enables changing of rounding mode for float-to-integer conversions, resulting in faster float-to-integer conversions.

### -size_lp64 (Itanium®-based systems only)

Default: Off

Specifies that 64-bit size for long and pointer types should be assumed.

### -[no]stack_temps

Default: Off

Specifies that arrays might be allocated on the stack, where possible, by the compiler.

### -nostartfiles

Default: Off

Specifies that standard startup files should be used when linking.

There is no `-startfiles` option.

### -syntax_only

Default: Off

Alternate syntax: `-y` and `-syntax`

Requests that only the syntax of the source file be checked. Code generation is suppressed.

### -T file

Default: Off

Instructs the linker to read link commands from `file`.

## -Tf file

Default: Off

Specifies that `filename` should be compiled as a Fortran source file. This option is used when you have a Fortran file with a nonstandard file extension (that is, not one of `.F`, `.FOR`, or `.F90`).

## -u

Default: Off

Alternate syntax: `-implicitnone`

Specifies that the IMPLICIT NONE should be set by default. See also `-warn [no]declarations`.

## -v

Default: Off

Specifies that driver tool commands should be shown and executed. See also `-dryrun`.

## -V

Default: None.

Displays the compiler version information.

## -what

Default: Off

Prints the version strings of the Fortran command and the compiler.

## -Wl,option1[,option2,...]

Default: Off

Passes options (specified by `option1`, `option2`, and so forth) to the linker for processing.

**-X**

Default: Off

Alternate syntax: `-nostdinc`

Removes standard directories from the include file search. This option prevents the compiler from searching the default path specified by the `FPATH` environment variable.

**-Xlinker value**

Default: Off

Passes `value` directly to the linker for processing.

# Optimization Options

The optimization options let you specify how to optimize your applications for speed, particular processors, code size, and so forth.

For more information about optimization, see "Compiler Optimizations Overview" and related sections in the *Intel Fortran User's Guide for Linux Volume II: Optimizing Applications*.

See also Floating-Point Options.

## Descriptions of Optimization Options

### -assume [no]buffered_io

Default: `-assume nobuffered_io` (buffer is flushed as each record is written)

Specifies whether records are written (flushed) to disk as each is written or are accumulated in the buffer. If you specify `-assume buffered_io`, records accumulate in the buffer.

For disk devices, `-assume buffered_io` (or the equivalent OPEN statement BUFFERED='YES' specifier or the `FORT_BUFFERED` run-time environment variable) requests that the internal buffer will be filled, possibly by many record output statements (WRITE), before it is written to disk by the Fortran run-time system. If a file is opened for direct access, I/O buffering will be ignored.

Using buffered writes usually makes disk I/O more efficient by writing larger blocks of data to the disk less often. However, if you request buffered writes, records not yet written to disk may be lost in the event of a system failure.

Unless you set the `FORT_BUFFERED` environment variable to true, the default is BUFFERED='NO' and `-assume nobuffered_io` for all I/O, in which case the Fortran run-time system empties its internal buffer for each WRITE (or similar record output statements).

The OPEN statement BUFFERED specifier applies to a specific logical unit. In contrast, the `-assume [no]buffered_io` option and the `FORT_BUFFERED` environment variable apply to all Fortran units.

## -auto_ilp32 (Itanium-based systems only)

Default: Off

Allows the compiler to use 32-bit pointers whenever possible as long as the application does not exceed a 32-bit address space.

Because this optimization requires interprocedural analysis over the whole program, you must use this option with the `-ipo` option.

Using this option on programs that exceed 32-bit address space may cause unpredictable results during program execution.

## -ax{K|W|N|B|P} (IA-32 systems only)

Default: None.

Directs the compiler to find opportunities to generate separate versions of functions that take advantage of features that are specific to the specified Intel® processor.

If the compiler finds such an opportunity, it first checks whether generating a processor-specific version of a function is likely to result in a performance gain. If this is the case, the compiler generates both a processor-specific version of a function and a generic version of the function. The generic version will run on any IA-32 processor.

At run time, one of the versions is chosen to execute, depending on the Intel processor in use. In this way, the program can benefit from performance gains on more advanced Intel processors, while still working properly on older IA-32 processors.

Possible values and the processors the code is optimized for are:

- `-axK`  Intel Pentium® III and compatible Intel processors
- `-axW`  Intel Pentium 4 and compatible Intel processors
- `-axN`  Intel Pentium 4 and compatible Intel processors
- `-axB`  Intel Pentium M and compatible Intel processors
- `-axP`  Intel® Pentium® 4 processors with Streaming SIMD Extensions 3 (SSE3) instruction support

## -complex_limited_range[-]

Default: Off (`-complex_limited_range-`)

Enables the use of basic algebraic expansions of some arithmetic operations involving data of type COMPLEX. This can result in performance improvements in programs that use a lot of COMPLEX arithmetic. However, values at the extremes of the exponent range might not compute correctly.

## -f[no-]alias

Default: `-falias`

Specifies that aliasing should be assumed in the program.

See also `-f[no-]fnalias`.

## -f[no-]fnalias

Default: `-ffnalias`

Specifies that aliasing should be assumed within functions. The `-fno-fnalias` option specifies that aliasiing should not be assumed within functions, but should be assumed across calls.

See also `-f[no-]alias`.

## -fast

Default: Off

Provides a shortcut method to enable several optimizations for run-time performance.

The `-fast` option sets the following options to improve performance:

- `-O3` (optimizes for maximum speed and high-level optimizations)
- `-ipo` (enables interprocedural optimizations across files)
- `-static` (prevents linking with shared libraries)

To get the best possible performance, you might need to use the option in conjunction with an architecture-specific option such as `-xN`.

To override one of the options set by `-fast`, specify that option after the `-fast` option on the command line.

### Note

The several options set by the `-fast` option might change from release to release.

## -fnsplit[-] (Itanium®-based systems only)

Default: Off

Enables function splitting if `-prof_use`. is also enabled. (This option has no effect if `-prof_use`. is not enabled.)

This option is automatically enabled if you use `-prof_use`.

To turn off function splitting, use `-fnsplit-`. (However, function grouping will continue to be enabled.)

See also these topics in *Volume II*:
Basic PGO Options
Example of Profile-Guided Optimization

## -fp (IA-32 systems only)

Default: On

Disables the use of ebp as a general-purpose register.

Most debuggers expect ebp to be used as a stack frame pointer, and cannot produce a stack backtrace unless this is so. This option allows frame pointers and disables the use of the ebp register in optimizations and lets the debugger produce a stack backtrace.

## -gp

Default: Off

Alternate syntax: `-p`

Compile and link for function profiling with the `gprof` tool.

## -ip

Default: Off

Enables single-file interprocedural optimizations.

Enhances inline function expansion.

See also this topic in Volume II: "Using -ip with -option Specifiers."

## -ip_no_inlining

Default: Off

Disables interprocedural inlining that results from the `-ip` or `-ipo`
interprocedural optimizations, but has no effect on other interprocedural
optimizations. Requires `-ip` or `-ipo`.

## -ip_no_pinlining (IA-32 systems only)

Default: Off

Disables partial inlining. Requires `-ip` or `-ipo`.

## -ipo

Default: Off

Enables Whole Program Optimization (WPO), which is the same as multifile
interprocedural optimization, or multifile IPO. All objects over the entire program
are compiled.

See also these topics in *Volume II*:
Multifile IPO Overview
Creating a Multifile IPO Executable with xilink
Using -Qip with -Qoption Specifiers

## -ipo_c

Default: Off

Optimizes across files and produces a multifile object file. Stops prior to the final link stage, leaving an optimized object file.

See also this topic in *Volume II*: "Analyzing the Effects of Multifile IPO."

## -ipo_obj

Default: Off

Forces the generation of real object files. Requires `-ipo.` See also this topic in *Volume II*: "Compilation with Real Object Files."

## -ipo_S

Default: Off

Optimizes across files and produces a multifile assembly file. Performs the same optimizations as `-ipo`, but stops prior to the final link stage, leaving an optimized assembly file. The default listing name is `ipo_out.s`.

See also this topic in *Volume II*: "Analyzing the Effects of Multifile IPO."

## -ivdep_parallel (Itanium®-based systems only)

Default: Off

Specifies that there is no loop-carried memory dependency in the loop where an IVDEP directive is specified. This technique is useful for some sparse matrix applications.

See also this topic in *Volume II*: "Memory Dependency with the IVDEP Directive."

## -nolib_inline

Default: On

Disables inline expansion of intrinsic functions.

## -On

Default: `-O2` unless you specify `-debug`, in which case the default is `-O0`

Specifies the code optimization for application types. Possible values are:

- `-O0`
  Disables all optimizations.
  This is the default if you specify `-debug` (with no keyword).
  Specifying this option causes certain `-warn` options to be ignored.
- `-O1`
  Alternate syntax on IA-32 systems: `-O2` or `-O`
  Maximize speed; disables some optimizations that increase code size for a small speed benefit. This option enables global optimization. This includes data-flow analysis, code motion, strength reduction and test replacement, split-lifetime analysis, and instruction scheduling. Specifying `-O2` includes the optimizations performed by `-O1`.
  Note that, on IA-32 systems, `-O1` and `-O2` are equivalent.
- `-O2`
  Alternate syntax on Itanium-based systems: `-O`
  Minimizes size; optimizes for speed, but disables some optimizations that increase code size for a small speed benefit; for the Itanium® compiler, `-O1` turns off software pipelining to reduce code size. This option enables local optimizations within the source program unit, recognition of common subexpressions, and expansion of integer multiplication and division using shifts.
- `-O3`
  Maximize speed plus use higher-level optimizations; optimizations include loop transformation, software pipelining, and (IA-32 only) prefetching; this option may not improve performance for some programs. Specifying `-O3` includes the optimizations performed by `-O2`. This option enables additional global optimizations that improve speed (at the cost of extra code size). These optimizations include:
  o Loop unrolling, including instruction scheduling
  o Code replication to eliminate branches
  o Padding the size of certain power-of-two arrays to allow more efficient cache use. (See also this topic in *Volume II*: "Using Arrays Efficiently.")
  Setting `-O3` sets `-fp`.

On IA-32 systems, `-O1`, `-O2`, and `-O` are equivalent.

On Itanium-based systems, `-O2` and `-O` are equivalent.

### Note

The last `-On` option specified on the command line takes precedence over any others.

## -opt_report

Default: Off

Generates an optimization report to `stderr`.

See also this topic in *Volume II*: "Optimizer Report Generation."

### -opt_report_file file

Default: Off

Generates an optimization report and specifies the file name for the report. You do not need to specify `-opt_report` if you use this option.

See also this topic in *Volume II*: "Optimizer Report Generation."

### -opt_report_help

Default: Off

Displays the optimization phases available for reporting.

See also this topic in *Volume II*: "Optimizer Report Generation."

### -opt_report_level {min|med|max}

Default: `-opt_report_level min`

Specifies the detail level of the optimization report.

See also this topic in *Volume II*: "Optimizer Report Generation."

### -opt_report_phase phase

Default: Off

Specifies the optimization phase to generate the report for. Can be specified multiple times on the command line for multiple optimizations.

See also this topic in *Volume II*: "Optimizer Report Generation."

### -opt_report_routine [routine]

Default: Off

Generates reports from all routines with names containing $routine$ as part of their name.

If the optional `routine` is not specified, reports from all routines are generated.

See also this topic in *Volume II*: "Optimizer Report Generation."

## -par_threshold n

Default: **-**`par_threshold 100`

Sets a threshold for the auto-parallelization of loops based on the probability of profitable parallel execution. `n` can be from 0 through 100.

`n=0`: loops get auto-parallelized regardless of computation work volume, that is, always.

`n=100`: loops get auto-parallelized only if profitable parallel execution is almost certain.

See also these topics in *Volume II*:

Auto-Parallelization Threshold Control and Diagnostics
Auto-Parallelization Overview
Auto-Parallelization: Enabling, Options, Directives, and Environment Variables

## -parallel

Default: Off

Enables the auto-parallelizer to generate multithreaded code for loops that can be safely executed in parallel. To use this option, you must also specify −O2 or −O3.

See also these topics in Volume II:
Auto-Parallelization Overview
Auto-Parallelization: Enabling, Options, Directives, and Environment Variables

## -prefetch[-] (IA-32 systems only)

Default: **-**`prefetch` (on)

Enables prefetch insertion optimization. The goal of prefetching is to reduce cache misses by providing hints to the processor about when data should be loaded into the cache. Note that −O3 must be specified for this option to work.

To disable the prefetch insertion optimization, use −`prefetch-`.

## -prof_dir dir

Default: The directory where the program is compiled.

Specifies the directory in which you intend to place the profiling output files (`.dyn` and `.dpi`) to be created. The specified directory must already exist.

See also these topics in *Volume II*:
Advanced PGO Options
Specific Coding Guidelines for IA-32 Architecture

## -prof_file file

Default: Source file name with extension `.dyn` and `.dpi`

Specifies the file name for the profiling summary file.

See also these topics in *Volume II*:
Advanced PGO Options
Specific Coding Guidelines for IA-32 Architecture

## -prof_gen

Default: Off

Instruments a program for profiling to get the execution count of each basic block.

See also these topics in Volume II:
Basic PGO Options
Example of Profile-Guided Optimization

## -prof_use

Default: Off

Enables use of profiling information (including function splitting and function grouping) during optimization. Instructs the compiler to produce a profile-optimized executable and merges available profiling output files into a `pgopti.dpi` file.

If you use this option, it automatically enables `-fnsplit[-]`.

Note that there is no way to turn off function grouping if you enable it using this option.

See also these topics in *Volume II*:
Basic PGO Options
Example of Profile-Guided Optimization

## -scalar_rep[-] (IA-32 systems only)

Default: `-scalar_rep` (on)

Enables scalar replacement performed during loop transformation. Requires `-O3`.

## -tpp*n*

Default value for IA-32 systems: `-tpp7`

Default value for Itanium®-based systems: `-tpp2`

Optimizes for a particular Intel® processor. The executable will run on other processors, but is optimized for processors noted below. Possible choices for $n$ are:

- 1   Optimize for Itanium processors (Itanium®-based systems only)
- 2   Optimize for Itanium 2 processors (Itanium®-based systems only)
- 5   Optimize for Intel Pentium® and Pentium® with MMX™ technology processors (IA-32 systems only)
- 6    Optimize for Intel Pentium® Pro, Pentium® II and Pentium® III processors (IA-32 systems only)
- 7    Optimize for Intel Pentium® 4, Intel® Xeon™, Intel Pentium® M processors, and Intel® Pentium® 4 processors with Streaming SIMD Extensions 3 (SSE3) instruction support (IA-32 systems only)

## -unroll[n]

Default: `-unroll` (lets the compiler decide)

Specifies the maximum number of times to unroll a loop.

Possible values are:

- `-unroll`   Lets the compiler decide.
- `-unroll0`   Disables loop unrolling. (Note: This is the only value allowed on Itanium-based systems; all other values are ignored.)
- `-unroll`*n*   Sets $n$ as the maximum number of times a loop can be unrolled.

## -x{K|W|N|B|P} (IA-32 systems only)

Intel® Fortran Compiler for Linux* Systems User's Guide Volume I: Building Applications

Default: None.

Lets you target your program to run on a specific Intel processor. The resulting code might contain unconditional use of features that are not supported on other processors.

Possible values and the processors the code is optimized for are:

- `-xK`  Intel Pentium III and compatible Intel processors
- `-xW`  Intel Pentium 4 and compatible Intel processors
- `-xN`  Intel Pentium 4 and compatible Intel processors
- `-xB`  Intel Pentium M and compatible Intel processors
- `-xP`  Intel Pentium 4 processors with Streaming SIMD Extensions 3 (SSE3) instruction support

To execute the program on x86 processors not provided by Intel Corporation, do not specify this option.

⚠️**Caution**

If a program compiled with this option is executed on a processor that lacks the specified set of instructions, it can fail with an illegal instruction exception, or display other unexpected behavior. In particular, programs compiled with `-xN`, `-xB`, or `-xP` will emit run-time errors if they are executed on unsupported processors.

# Output Files Options

The output options let you specify names and directory locations for files that result from the compilation.

## Descriptions of Output Files Options

**-c**

Default: Off

Specifies that the compiler should compile to object (`.o` file) only and not link.

**-fcode-asm**

Default: Off

Specifies that the compiler should produce an assembly file with optional code annotations. This option requires the use of the `-S` option.

### -fsource-asm

Default: Off

Specifies that the compiler should produce an assembly file with optional source annotations. This option requires the use of the `-S` option.

### -f[no]verbose-asm

Default: On when `-S` is specified

Specifies that the compiler should produce an assembly file with compiler comments, including options and version information.This option requires the use of the `-S` option.

### -module path

Specifies the directory (*path*) where module files (file extension `.mod`) are placed. See also Searching for Include and .mod Files,

### -ofilename

Default: Off

Specifies the name of the output file.

If `-c` is specified, `-o` specifies the name of the object file.

If `-S` is specified, `-o` specifies the name of the assembly listing file.

If neither `-c` nor `-S` is specified, `-o` specifies the name of the executable file.

### -Qinstall dir

Default: Off

Specifies *dir* as the root directory for the compiler installation.

### -Qlocation,tool,path

Default: Off

Specifies the directory location of supporting tools, specifically the preprocessor, compiler, assembler, and linker.

For syntax and details, see Using -Qlocation to Specify an Alternate Location for a Tool.

### -Qoption,tool,options

Default: Off

Passes options to tools, specifically the preprocessor, compiler, assembler, and linker.

For syntax and details, see Using -Qoption to Pass Options to Tools.

### -S

Default: Off

Specifies that the compiler should compile to assembly (`.s`) file only and not link.

### -use_asm

Default: Off

Specifies that objects should be produced through the assembler.

# Preprocessor Options

The preprocessor options let you specify how the compiler preprocesses files as an optional first phase of the compilation and where it looks for source directories.

**See Also**
Preprocessing Phase
Predefined Preprocessor Symbols

## Descriptions of Preprocessor Options

### -assume [no]source_include

Default: `-assume source_include`

Specifies the directory searched for module files specified by a USE statement or source files specified by an INCLUDE statement.

Possible values are:

- `-assume source_include`   Search in source file directory
- `-assume nosource_include`  Search in current directory

Note that you can use this option whether or not you use the `-fpp` option.

## -Dname[=value]

Default: Off

Specifies one or more definitions for use with conditional compilation directives or the Fortran preprocessor, fpp. If you have more than one, use separate `-D` options.

The *value* can be a character or integer value. If *value* not specified, 1 is assumed to be the value for *name*.

For an example, see Defining Preprocessor Symbols.

### Note

Do not use `D` for *name*, because it will conflict with the `-DD` option (alternate syntax for `-d_lines`). However, you can use the `-Dname=n` syntax, such as `-DD=1`.

## -[no]fpp

Default: `-nofpp`

Alternate syntax: `-[no]cpp`; also `fpp 0` is equivalent to `-nofpp`; also `fpp n` (where *n* is any number greater than 0) is equivalent to `-fpp`

Invokes the FPP Preprocessor (fpp) prior to compilation, enabling preprocessor directives.

`-cpp` is the same as `-fpp` (runs fpp, not the C preprocessor).

## -Idir

Default: Include path

Specifies one or more directories to add to the include path, which is searched for module files (USE statement) and include files (INCLUDE statement). Use a semicolon delimiter for more than one directory.

To request that the compiler search first in the directory where the source file resides instead of the current directory, specify `-assume source_include`.

For all USE statements and for those INCLUDE statements whose file name does not begin with a device or directory name, the directories searched are as follows, in this order:

1. The directory containing the first source file (if `-assume source_include` was specified, which is the default).
2. The current default directory where the compilation is taking place.
3. If specified, the directory or directories listed in the `-Idir` option. The order of searching multiple directories occurs within the specified list from left to right
4. The directories indicated in the compile-time environment variable `FPATH`.

See also `-noinclude`.

## -preprocess_only

Default: Off

Alternate syntax: `-P` and `-F`

Runs the Fortran preprocessor (fpp) and puts the result for each source file in a corresponding `.i` or `.i90` file. The `.i` or `.i90` file does not have line numbers (#) in it. The source file is not compiled.

You need to specify `-fpp` with this option.

## -U name

Default: Off.

Undefines any definition currently in effect for the symbol specified by `name`.

## -Wp,option1[,option2,...]

Default: Off

Passes options (specified by `option1`, `option2`, and so forth) to the preprocessor.

This option is the same as `-fpp`, except that `-fpp` also invokes fpp.

# Run-Time Options

The run-time options let you specify error checking to be performed at run time, not compile time.

## Descriptions of Run-Time Options

### -[no]check [all] or -[no]check [none]

Default: Custom (options are specified individually).

Alternate syntax: `-C` is equivalent to `-check all`

Specifies all or no checking for run-time failures. Individual run-time check options shown below are not available if `-check all` or `-check none` is specified.
`-nocheck` is equivalent to `-check none`.

`-check` is equivalent to `-check all`.

### -check [no]arg_temp_created

Default: Custom (options are specified individually).

Default: `-check noarg_temp_created`

Requests a run-time informational message if actual arguments are copied into temporary storage before routine calls.

### -check [no]bounds

Default: `-check nobounds`

Alternate syntax: `-CB`

Generates code to perform run-time checks on array subscript and character substring expressions.

®

The default (`-check nobounds`) suppresses range checking.

For array bounds, each individual dimension is checked. Array bounds checking is not performed for arrays that are dummy arguments in which the last dimension bound is specified as * or when both upper and lower dimensions are 1.

Once the program is debugged, omit this option to reduce executable program size and slightly improve run-time performance.

## -check [no]format

Default: `-check noformat,` unless `-vms` is specified, in which case `-check format` is the default.

Requests a run-time error message when the data type for an item being formatted for output does not match the FORMAT descriptor.

## -check [no]output_conversion

Default: `-check nooutput_conversion,` unless `-vms` is specified, in which case `-check output_conversion` is the default

Requests a run-time error message when format truncation occurs (that is, when a number is too large to fit in the specified format field length without loss of significant digits).

## -[no]traceback

Default: `-notraceback`

Requests that the compiler generate extra information in the object file that allows the display of source file traceback information at run time when a severe error occurs.

Specifying `-traceback` provides source file, routine name, and line number correlation information in the displayed call stack hexadecimal addresses (program counter trace) that is displayed when a severe error occurs. If `-traceback` is not specified, this information is not displayed. However, advanced users can locate the cause of the error using a map file and the hexadecimal addresses of the stack displayed when a severe error occurs.

Specifying `-traceback` will increase the size of the executable program, but has no impact on run-time execution speeds.

The `-traceback` option functions independently of the `-debug` option.

**Debugging Using idb**

# Debugging Using idb Overview

See these topics:

Getting Started with Debugging

Preparing Your Program for Debugging

Using Debugger Commands and Setting Breakpoints

Summary of Debugger Commands

Debugging the SQUARES Example Program

Displaying Variables

Expressions in Debugger Commands

Debugging Mixed-Language Programs

Debugging a Program that Generates a Signal

Locating Unaligned Data

# Getting Started with Debugging

The Intel® Debugger (idb) is a source-level, symbolic debugger that lets you:

- Control the execution of individual source lines in a program.
- Set stops (breakpoints) at specific source lines or under various conditions.
- Change the value of variables in your program.
- Refer to program locations by their symbolic names, using the debugger's knowledge of the Intel Fortran language to determine the proper scoping rules and how the values should be evaluated and displayed.
- Print the values of variables and set a tracepoint (trace) to notify you when the value of a variable changes. (Another term for a tracepoint is a watchpoint.)
- Perform other functions, such as examining core files, examining the call stack, or displaying registers.

Intel® Fortran Compiler for Linux* Systems User's Guide Volume I: Building Applications

The idb debugger has two modes:

- dbx (default mode)
- gdb (optional mode)

All examples in this guide are shown in dbx mode.

**Note**

For complete information about idb, see the `idb` man page or the online *Intel® Debugger (IDB) Manual.*

## Debugging Options

To use the debugger, you should specify the `ifort` command and the `-g` command-line option. Traceback information and symbol table information are both necessary for debugging. If you specify `-g`, the compiler provides the symbol table and traceback information needed for symbolic debugging. (The `-notraceback` option cancels the traceback information.)

Likely uses of these options at the various stages of program development are as follows:

During early stages of program development, use the `-g` option to create unoptimized code (optimization level `-O0`). This option also might be chosen later to debug reported problems from later stages.

Traceback and symbol table information result in a larger object file. During the later stages of program development, use `-g0` or `-g1` to minimize the object file size and, as a result, the memory needed for program execution, usually with optimized code. (The `-g0` option eliminates the traceback information.)

When you have finished debugging your program, you can recompile and relink to create an optimized executable program or remove traceback and symbol table information with the `strip` command. (See `strip(1)`.)

Using the `-help` option to the compiler indicates `-g` only.

Note: Debugging of optimized code is not fully supported on Intel platforms.

# Preparing Your Program for Debugging

Use the `ifort` command with certain options to create an executable program
for debugging. To invoke the debugger, enter the debugger shell command and
the name of the executable program.

The following commands create (compile and link) the executable program and
invoke the interface to the debugger:

```
ifort -g -o squares squares.f90
idb squares
Linux Application Debugger for xx-bit applications, Version
x.x, Build xxxx
object file name: squares
reading symbolic information ... done
(idb)
```

In this example, the `ifort` command:

- Compiles and links the program `squares.f90`.
- Requests symbol table information needed for symbolic debugging and no
  optimization (`-g`).
- Names the executable file `squares` instead of `a.out` (`-o squares`).

The `idb` shell command runs the debugger, specifying the executable program
`squares`.

At the debugger prompt `(idb)`, you can enter a debugger command.

See also the online *Intel ® Debugger (IDB) Manual.*

# Using Debugger Commands and Setting Breakpoints

To find out what happens at critical points in your program, you need to stop
execution at these points and look at the contents of program variables to see if
they contain the correct values. Points at which the debugger stops program
execution are called breakpoints.

To set a breakpoint, use one of the forms of the `stop` or `stopi` commands.

Using a sample program, the following debugger commands set a breakpoint at
line 4, run the program, continue the program, delete the breakpoint, rerun the
program, and return to the shell:

```
(idb) stop at 4
```

```
[#1: stop at "squares.f90":4 ]
(idb) run
[1] stopped at [squares:4 0x120001880]
>   4       OPEN(UNIT=8, FILE='datafile.dat', STATUS='OLD')
(idb) cont
Process has exited with status 0
(idb) delete 1
(idb) rerun
Process has exited with status 0
(idb) quit
%
```

In this example:

- The `stop at 4` command sets a breakpoint at line 4. To set a breakpoint at the start of a subprogram (such as `calc`), use the `stop in` command (such as `stop in calc`).
- The `run` command begins program execution and stops at the first breakpoint. The program is now active, allowing you to view the values of variables with print commands and perform related functions.
- The `cont` command resumes (continues) program execution. In addition to the `cont` command, you can also use the `step`, `next`, `run`, or `rerun` commands to resume execution.
- The `delete 1` command shows how to delete a previously set breakpoint (with event number 1). For instance, you might need to delete a previously set breakpoint before you use the `rerun` command.
- The `rerun` command runs the program again. Since there are no breakpoints, the program runs to completion.
- The `quit` command exits the debugger and returns to the shell.

# Other Debugger Commands

Other debugger commands include the following:

- To get help on debugger commands, enter the `help` command.
- To display previously typed debugger commands, enter the `history` command.
- To examine the contents of a location, use the `print` or `dump` commands.
- To execute a shell command, use the `sh` command (followed by the desired shell command). For instance, if you cannot recall the name of a FUNCTION statement, the following `grep` shell command displays the lines containing the letters FUNCTION, allowing use of the function name (SUBSORT) in the stop in command:

```
(idb) sh grep FUNCTION data.for
INTEGER*4 FUNCTION SUBSORT (A,B)
(idb) stop in subsort
(idb)
```

See also Summary of Debugger Commands.

# Summary of Debugger Commands

The table below lists some of the more frequently used debugging commands available in idb.  Many of these commands can be abbreviated (for example, you can enter `c` instead of `cont` and `s` instead of `step`.  You can use the `alias` command to get a complete list of these abbreviations and even create your own aliases.

The table below shows examples of the most commonly used debugger commands. For more information, see the online *Intel ® Debugger (IDB) Manual.*

| Command example | Description |
|---|---|
| `catch` | Displays all signals that the debugger is currently set to catch. See also `ignore`. |
| `catch fpe` | Tells the debugger to catch the fpe signal (or the signal specified). This prevents the specified signal from reaching the Intel Fortran run-time library (RTL). |
| `catch unaligned` | Tells the debugger to catch the unaligned signal. |
| `cont` | Resumes (continues) execution of the program that is being debugged. Note that there is no idb command named `continue`. |
| `delete 2` | Removes the breakpoint or tracepoint identified by event number 2. See also `status`. |
| `delete all` | Removes all breakpoints and tracepoints. |
| `help` | Displays debugger help text. |
| `history 5` | Displays the last five debugger commands. |
| `ignore` | Displays the signals the debugger is currently set to ignore. The ignored signals are allowed to pass directly to the Intel Fortran RTL, See also `catch`. |
| `ignore fpe` | Tells the debugger to ignore the fpe signal (or the signal specified). This allows the specified signal to pass directly to the Intel Fortran RTL, allowing message display. |
| `ignore unaligned` | Tells the debugger to ignore the unaligned signal (the default). |
| `kill` | Terminates the program process, leaving the debugger |

| | running and its breakpoints and tracepoints intact for when the program is rerun. |
|---|---|
| `list` | Displays source program lines. To list a range of lines, add the starting line number, a comma (,), and the ending line number, such as `list 1,9`. |
| `print k` | Displays the value of the specified variable, such as `K`. |
| `printregs` | Displays all registers and their contents. |
| `next` | Steps one source statement but does *not* step into calls to subprograms, Compare with `step`. |
| `quit` | Ends the debugging session. |
| `run` | Runs the program being debugged. You can specify program arguments and redirection. |
| `rerun` | Runs the program being debugged again. You can specify program arguments and redirection. |
| `return [routine-name]` | Continues execution of the function until it returns to its caller.<br><br>When using the `step` command, if you step into a subprogram that does not require further investigation, use the `return` command to continue execution of the current function until it returns to its caller. If you include the name of a routine with the `return` command, execution continues until control is returned to that routine.<br><br>The *routine-name* is the name of the routine, usually named by a PROGRAM, SUBROUTINE, or FUNCTION statement. If there is no PROGRAM statement, the debugger refers to the main program with a prefix of `main$` followed by the file name. |
| `sh more progout.f90` | Executes the shell command `more` to display file `progout.f90`, then returns to the debugger environment. |
| `show thread` | Lists all threads known to the  debugger. |
| `status` | Displays breakpoints and tracepoints with their event numbers. See also `delete`. |
| `step` | Steps one source statement, including stepping *into* calls of a subprogram. For Intel Fortran I/O statements, intrinsic procedures, library routines, or other subprograms, use the `next` command instead of `step` to step over the subprogram call. Compare with `next`; see also  `return`. |
| `stop in foo` | Stops execution (breakpoint) at the beginning of routine `foo`. |
| `stop at 100` | Stops execution at line 100 (breakpoint) of the current source file. |
| `stopi at xxxxxxx` | Stops execution at address *xxxxxxx* of the current executable program. |
| `thread [n]` | Identifies or sets the current thread context. |

| watch *location* | Displays a message when the debuggee (or user program) accesses the specified memory location. For example: `watch 0x140000170` |
|---|---|
| watch variable *m* | Displays a message when the debuggee (or user program) accesses the variable specified by *m*. |
| whatis *symbol* | Displays the data type of the specified symbol. |
| when at *9* {*command*} | Executes a command or commands.<br><br>When a specified line (such as 9) is reached, the *command* or commands are executed. For example, `when at 9 {print k}` prints the value of variable `K` when the program executes source code line 9. |
| when in *name* {*command*} | Executes a command or commands.<br><br>When a procedure specified by *name* is reached, the *command* or commands are executed. For example, `when in calc_ave {print k}` prints the value of variable `K` when the program begins executing the procedure named `calc_ave`. |
| where | Displays the call stack. |
| where thread all | Displays the stack traces of all threads. |

The debugger supports other special-purpose commands. For example:

- You might use the `attach` and `detach` commands for programs with very long execution times.
- The `listobj` command might be helpful when debugging programs that depend on shared libraries. The `listobj` command displays the names of executables and shared libraries currently known to the debugger.

# Debugging the SQUARES Example Program

The example below shows a program called SQUARES that requires debugging. The program was compiled and linked without diagnostic messages from either the compiler or the linker. However, this program contains a logic error in an arithmetic expression.

Compiler-assigned line numbers have been added in the example so that you can identify the source lines to which the explanatory text refers.

```
PROGRAM SQUARES
    INTEGER INARR(10), OUTARR(10), I, K
```

```
! Read the input array from the data file.
      OPEN(UNIT=8, FILE='datafile.dat', STATUS='OLD')
      READ(8,*,END=5) N, (INARR(I), I=1,N)
  5   CLOSE (UNIT=8)

! Square all nonzero elements and store in OUTARR.
      K = 0
      DO I = 1, N
        IF (INARR(I) .NE. 0) THEN
        K = K + 1         ! add this line
        OUTARR(K) = INARR(I)**2
        ENDIF
      END DO

! Print the squared output values.  Then stop.
      PRINT 20, N
  20  FORMAT (' Total number of elements read is',I4)
      PRINT 30, K
   0  FORMAT (' Number of nonzero elements is',I4)
      DO, I=1,K
         PRINT 40, I, OUTARR(K)
  40     FORMAT(' Element', I4, 'Has value',I6)
      END DO
   END PROGRAM SQUARES
```

The program SQUARES performs the following functions:

1.  Reads a sequence of integer numbers from a data file and saves these numbers in the array INARR. The file datafile.dat contains one record with the integer values 4, 3, 2, 5, and 2. The first number indicates the number of data items that follow.
2.  Enters a loop in which it copies the square of each nonzero integer into another array OUTARR.
3.  Prints the number of nonzero elements in the original sequence and the square of each such element.

Note: This example assumes that the program was executed without array bounds checking (set by the -check bounds command-line option). When executed with array bounds checking, a run-time error message appears.

When you run SQUARES, it produces the following output, regardless of the number of nonzero elements in the data file:

```
squares
Number of nonzero elements is   0
```

The logic error occurs because variable K, which keeps track of the current index into OUTARR, is not incremented in the loop on lines 9 through 13. The statement K = K + 1 should be inserted just before line 11.

The following example shows how to start the debugging session and how to use the character-cell interface to idb to find the error in the sample program shown earlier. Comments keyed to the callouts at the right follow the example:

```
ifort -g -o squares squares.f90   (1)
idb squares                       (2)
Linux Application Debugger for xx-bit applications, Version
x.x, Build xxxx
object file name: squares
reading symbolic information ... done
(idb) list 1,9                            (3)
     1    PROGRAM SQUARES
     2      INTEGER INARR(20), OUTARR(20)
     3 C  ! Read the input array from the data file.
>    4       OPEN(UNIT=8, FILE='datafile.dat', STATUS='OLD')
     5       READ(8,*,END=5) N, (INARR(I), I=1,N)
     6    5 CLOSE (UNIT=8)
     7 C  ! Square all nonzero elements and store in
OUTARR.
     8       K = 0
     9        DO 10 I = 1, N
(idb) stop at 8                           (4)
[#1: stop at "squares.f90":8]
(idb) run                                 (5)
[1] stopped at ["squares.f90":4 0x120001a88]
>    8       K = 0
(idb) step                                (6)
stopped at [squares:9 0x120001a90]
     9         DO 10 I = 1, N
(idb) print n, k                          (7)
4 0
(idb) step                                (8)
stopped at [squares:10 0x120001ab0]]
    10        IF(INARR(I) .NE. 0)  THEN
(idb) s
stopped at [squares:11 0x1200011acc]
    11        OUTARR(K) =   INARR(I)**2
(idb) print i, k                          (9)
1 0
(idb) assign k = 1                        (10)
(idb) watch variable k                    (11)
[#2: watch variable (write) k 0x1400002c0 to 0x1400002c3 ]
(idb) cont                                (12)
Number of nonzero elements is   1
Element   1 has value     4
```

```
Process has exited with status 0
(idb) quit                                      (13)
% vi squares.f90                                     (14)
   .
   .
   .
10:        IF(INARR(I) .NE. 0) THEN
11:           K = K + 1
12:           OUTARR(K) = INARR(I)**2
13:        ENDIF
   .
   .
   .
% ifort -g -o squares squares.f90       (15)
% idb squares
Welcome to the idb Debugger Version x.x-xx
Reading symbolic information ...done
(idb) when at 12 {print k}                  (16)
[#1: when at "squares.f90":12 { print K} ]
(idb) run                                       (17)
[1] when [squares:12 0x120001ae0]
1
[1] when [squares:12 0x120001ae0]
2
[1] when [squares:12 0x120001ae0]
3
[1] when [squares:12 0x120001ae0]
4
Number of nonzero elements is    4
Element    1 has value       9
Element    2 has value       4
Element    3 has value      25
Element    4 has value       4
Process has exited with status 0
(idb) quit                                       (18)
%
```

1. On the command line, the `-g` option directs the compiler to write the symbol information associated with SQUARES into the object file for the debugger. It also disables most optimizations done by the compiler to ensure that the executable code matches the source code of the program.
2. The shell command `idb squares` runs the debugger, which displays its banner and the debugger prompt, `(idb)`. This command specifies the executable program as a file named `squares`. You can now enter debugger commands. After the `idb squares` command, execution is initially paused before the start of the main program unit (before program SQUARES, in this example).
3. The `list 1,9` command prints lines 1 through 9.
4. The command `stop at 8` sets a breakpoint (1) at line 8.

5. The `run` command begins program execution. The program stops at the first breakpoint, line 8, allowing you to examine variables `N` and `K` before program completion.
6. The `step` advances the program to line 9. The `step` command ignores source lines that do not result in executable code; also, by default, the debugger identifies the source line at which execution is paused. To avoid stepping into a subprogram, use the `next` command instead of `step`.
7. The command `print n, k` displays the current values of variables N and K. Their values are correct at this point in the execution of the program.
8. The two `step` commands continue executing the program into the loop (lines 9 to 11) that copies and squares all nonzero elements of INARR into OUTARR. Certain commands can be abbreviated. In this example, the s command is an abbreviation of the `step` command.
9. The command `print i, k` displays the current values of variables `I` and `K`. Variable `I` has the expected value, 1. But variable `K` has the value 0 instead of the expected value, 1. To fix this error, `K` should be incremented in the loop just before it is used in line 11.
10. The `assign` command assigns `K` the value 1.
11. The `watch variable k` command sets a watchpoint that is triggered every time the value of variable `K` changes. In the original version of the program, this watchpoint is never triggered, indicating that the value of variable `K` never changes (a programming error).
12. To test the patch, the `cont` command (an abbreviation of `continue`) resumes execution from the current location. The program output shows that the patched program works properly, but only for the first array element. Because the watchpoint (`watch variable k` command) does not occur, the value of `K` did not change and there is a problem. The idb message "Process has exited with status 0" shows that the program executed to completion.
13. The `quit` command returns control to the shell so that you can correct the source file and recompile and relink the program.
14. The shell command `vi` runs a text editor and the source file is edited to add K = K + 1 after line 10, as shown. (Compiler-assigned line numbers have been added to clarify the example.)
15. The revised program is compiled and linked. The shell command `idb squares` starts the debugger, using the revised program so that its correct execution can be verified.
16. The `when at 12 {print k}` command reports the value of `K` at each iteration through the loop.
17. The `run` command starts execution.
18. The displayed values of `K` confirm that the program is running correctly.
19. The `quit` command ends the debugging session, returning control to the shell.

# Displaying Variables in the Debugger

To refer to a variable, use either the uppercase or lowercase letters. For example:

```
(idb) print J
(idb) print j
```

You can enter command names in uppercase:

```
(idb) print J
```

If you compile the program with the command-line option `-names as_is` and you need to examine case-sensitive names, you can control whether idb is case-sensitive by setting the `$lang` environment variable to the name of a case-sensitive language.

## Module Variables

To refer to a variable defined in a module, insert an opening quote ('), the module name, and another opening quote (') before the variable name. For example, with a variable named `J` defined in a module named `modfile` (statement MODULE MODFILE), enter the following command to display its value:

```
(idb)     list 5,9
      5      USE MODFILE
      6       INTEGER*4 J
      7       CHARACTER*1 CHR
      8       J = 2**8
(idb)     print 'MODFILE'J
256
```

## Common Block Variables

You can display the values of variables in a Fortran common block by using debugger commands such as `print` or `whatis`.

To display the entire common block, use the common block name.

To display a specific variable in a common block, use the variable name. For example:

```
(idb)     list 1,11
   1       PROGRAM EXAMPLE
   2
```

```
   3          INTEGER*4 INT4
   4          CHARACTER*1 CHR
   5          COMMON /COM_STRA/ INT4, CHR
   6
   7          CHR = 'L'
   8
   9          END
(idb)      print COM_STRA
COMMON
        INT4 = 0
        CHR = "L"
(idb)
(idb)      print CHR
"L"
```

If the name of a data item in a common block has the same name as the common block itself, the data item is accessed.

# Derived-Type Variables

Variables in a Fortran 95/90 derived-type (TYPE statement) are represented in idb commands such as `print` or `whatis` using Fortran 95/90 syntax form.

For derived-type structures, use the derived-type variable name, a percent sign (%), and the member name. For example:

```
(idb)      list 3,11
   3  TYPE X
   4      INTEGER A(5)
   5   END TYPE X
   6
   7  TYPE (X) Z
   8
   9  Z%A = 1
  10
  11  PRINT *,Z%A
(idb)      print Z%A
(1) 1
(2) 1
(3) 1
(4) 1
(5) 1
(idb)
```

To display the entire object, use the `print` command with the object name. For example:

```
(idb)  print Z
```

# Record Variables

To display the value of a field in a record structure, enter the variable name as: the record name, a delimiter (either a period (.) or a percent sign (%)), and the field name.

To view all fields in a record structure, enter the name of the record structure, such as REC (instead of REC.CHR or REC%CHR) in the previous example.

# Pointer Variables

Intel Fortran supports two types of pointers:

- Fortran 95/90 pointers (standard-conforming)
- Integer pointers (extension to the Fortran 95/90 standards)

### Fortran 95/90 Pointers

Fortran 95/90 pointers display their corresponding target data with a `print` command.

Comments keyed to the callouts at the right follow the example:

```
ifort -g point.f90
idb ./a.out
Linux Application Debugger for xx-bit applications, Version
x.x, Build xxxx
object file name: ./a.out
Reading symbolic information ...done
(idb) stop in ptr
[#1: stop in ptr ]
(idb) list 1:13
     1        program ptr
     2
     3        integer, target :: x(3)
     4        integer, pointer :: xp(:)
     5
     6        x = (/ 1, 2, 3/)
     7        xp => x
     8
     9        print *, "x = ", x
    10        print *, "xp = ", xp
    11
    12        end
(idb) run
[1] stopped at [ptr:6 0x120001838]
     6         x = (/ 1, 2, 3/)
```

```
(idb) whatis x
int x(1:3)
(idb) whatis xp                                          (1)
int xp(:)
(idb) s
stopped at [ptr:7 0x120001880]
      7         xp => x
(idb) s
stopped at [ptr:9 0x120001954]
      9         print *, "x = ", x
(idb) s
x =              1              2              3
stopped at [ptr:10 0x1200019c8]
(idb) s
 xp =             1             2             3
stopped at [point:12 0x120001ad8]
     12         end
(idb) S
 xp =             1             2             3
(idb) whatis xp                                          (2)
int xp(1:3)
(idb) print xp
(1) 1
(2) 2
(3) 3
(idb) quit
%
```

1. For the first `whatis xp` command, `xp` has not yet been assigned to point to variable `x` and is a generic pointer.
2. Since `xp` has been assigned to point to variable `x`, for the second `whatis xp` command, `xp` takes the same size, shape, and values as `x`.

## Integer Pointers

Like Fortran 95/90 pointers, integer pointers (also known as Cray*-style pointers) display the target data in their corresponding source form with a `print` command:

```
(idb) stop at 14
[#1: stop at "dfpoint.f90":14 ]
(idb) run
[1] stopped at [dfpoint:14 0x1200017e4]
(idb) list 1,14
      1         program dfpoint
      2
      3         real i(5)
      4         pointer (p,i)
      5
      6         n = 5
```

```
     7
     8          p = malloc(sizeof(i(1))*n)
     9
    10          do j = 1,5
    11              i(j) = 10*j
    12          end do
    13
>   14          end
(idb) whatis p
float (1:5) pointer p
(idb) print p
0x140003060 = (1) 10
(2) 20
(3) 30
(4) 40
(5) 50
(idb) quit
%
```

# Array Variables

For array variables, put subscripts within parentheses, as with Fortran 95/90 source statements. For example:

```
(idb) assign arrayc(1)=1
```

You can print out all elements of an array using its name. For example:

```
(idb) print arrayc
(1) 1
(2) 0
(3) 0
(idb)
```

Avoid displaying all elements of a large array. Instead, display specific array elements or array sections. For example, to print array element arrayc(2):

```
(idb) print arrayc(2)
(2) 0
```

## Array Sections

An array section is a portion of an array that is an array itself. An array section can use subscript triplet notation consisting of a three parts: a starting element, an ending element, and a stride.

Consider the following array declarations:

```
INTEGER, DIMENSION(0:99)     :: arr
INTEGER, DIMENSION(0:4,0:4)  :: FiveByFive
```

 Assume that each array has been initialized to have the value of the index in each position, for example, FiveByFive(4,4) = 44, arr(43) = 43. The following examples are array expressions that will be accepted by the debugger:

```
(idb) print arr(2)
2
(idb) print arr(0:9:2)
(0) = 0
(2) = 2
(4) = 4
(6) = 6
(8) = 8
(idb) print FiveByFive(:,3)
(0,3) = 3
(1,3) = 13
(2,3) = 23
(3,3) = 33
(4,3) = 43
(idb)
```

The only operations permissible on array sections are whatis and print.

### Assignment to Arrays

Assignment to array elements are supported by idb.

For information about assigning values to whole arrays and array sections, see the Fortran chapter in the online *Intel® Debugger (IDB) Manual.*

## Complex Variables

idb supports COMPLEX or COMPLEX*8, COMPLEX*16, and COMPLEX*32 variables and constants in expressions.

Consider the following Fortran program:

```
PROGRAM complextest
      COMPLEX*8 C8 /(2.0,8.0)/
      COMPLEX*16 C16 /(1.23,-4.56)/
      REAL*4    R4 /2.0/
      REAL*8    R8 /2.0/
      REAL*16   R16 /2.0/

      TYPE *, "C8=", C8
```

```
        TYPE *, "C16=", C16
     END PROGRAM
```

idb supports the display and assignment of COMPLEX variables and constants
as well as basic arithmetic operators. For example:

```
Welcome to the idb Debugger Version x.x-xx
------------------
object file name: complex
Reading symbolic information ...done
(idb) stop in complextest
[#1: stop in complextest ]
(idb) run
[1] stopped at [complextest:15 0x1200017b4]
     15        TYPE *, "C8=", C8
(idb) whatis c8
complex c8
(idb) whatis c16
double complex c16
(idb) print c8
(2, 8)
(idb) print c16
(1.23, -4.56)
(idb) assign c16=(-2.3E+10,4.5e-2)
(idb) print c16
(-23000000512, 0.04500000178813934)
(idb)
```

## Data Types

The table below shows the Intel Fortran data types and their equivalent built-in
debugger names:

| Fortran 95/90 data type declaration | Debugger equivalent |
|---|---|
| CHARACTER | character |
| INTEGER, INTEGER(KIND=*n*) | integer, integer*n |
| LOGICAL, LOGICAL (KIND=*n*) | logical, logical*n |
| REAL, REAL(KIND=4) | real |
| DOUBLE PRECISION, REAL(KIND=8) | real*8 |
| REAL(KIND=16) | real*16 |
| COMPLEX, COMPLEX(KIND=4) | complex |
| DOUBLE COMPLEX, COMPLEX(KIND=8) | double complex |
| COMPLEX(KIND=16), COMPLEX*32 | long double complex |

# Expressions in Debugger Commands

Expressions in debugger commands use Fortran 95/90 source language syntax for operators and expressions.

Enclose debugger command expressions between curly braces ({ }). For example, the expression `print k` in the following statement is enclosed between curly braces ({ }):

```
(idb) when at 12 {print k}
```

## Fortran Operators

The Intel Fortran operators include the following:

- Relational operators, such as less than (.LT. or <) and equal to (.EQ. or ==)
- Logical operators, such as logical conjunction (.AND.) and logical disjunction (.OR.)
- Arithmetic operators, including addition (+), subtraction (--), multiplication (*), and division (/).

For a complete list of operators, see the *Intel Fortran Language Reference Manual.*

## Procedures

The idb debugger supports invocation of user-defined specific procedures using Fortran 95/90 source language syntax.

**See Also**

*Intel Fortran Language Reference Manual*

Online *Intel ® Debugger (IDB) Manual*

# Debugging Mixed-Language Programs

The idb debugger lets you debug mixed-language programs. Program flow of control across subprograms written in different languages is transparent.

The debugger automatically identifies the language of the current subprogram or code segment on the basis of information embedded in the executable file. For

example, if program execution is suspended in a subprogram in Fortran, the current language is Fortran. If the debugger stops the program in a C function, the current language becomes C. The debugger uses the current language to determine the valid expression syntax and the semantics used to evaluate an expression.

The debugger sets the `$lang` environment variable to the language of the current subprogram or code segment. By manually setting the `$lang` environment variable, you can force the debugger to interpret expressions used in commands by the rules and semantics of a particular language. For example, you can check the current setting of `$lang` and change it as follows:

```
(idb) print $lang
"C++"
(idb) set $lang = "Fortran"
```

When the `$lang` environment variable is set to "Fortran", names are case-insensitive. To make names case-sensitive when the program was compiled with the `-names as_is` option, specify another language for the `$lang` environment variable, such as C, view the variable, then set the `$lang` environment variable to "Fortran".

# Debugging a Program that Generates a Signal

If your program encounters a signal (exception) at run time, to make it easier to debug the program, you should recompile and relink with the following command-line options before debugging the cause of the signal:

- Use the `-fpen` option to control the handling of signals.
- As with other debugging tasks, use the `-g` option to generate sufficient symbol table information and debug unoptimized code.

If requested, idb will catch and handle signals before the Intel Fortran run-time library (RTL) does. You can use the idb commands `catch` and `ignore` to control whether idb catches signals or ignores them:

- When idb catches a sig0nal, an idb message is displayed and execution stops at that statement line. The error-handling routines provided by the RTL are not called. At this point, you can examine variables and determine where in the program the signal has occurred.
- When idb ignores a signal, the signal is passed to the RTL. This allows the handling and display of run-time signal messages in the manner requested during compilation.

To obtain the appropriate run-time error message when debugging a program that generates a signal (especially one that allows program continuation), you might need to use the `ignore` command before running the program. For instance, use the following command to tell the debugger to ignore floating-point signals and pass them through to the RTL:

(idb) **ignore fpe**

In cases where you need to locate the part of the program causing a signal, consider using the `where` command.

# Locating Unaligned Data

Unaligned data can slow program execution. You should determine the cause of the unaligned data, fix the source code (if necessary), and recompile and relink the program.

If your program encounters unaligned data at run time, to make it easier to debug the program, you should recompile and relink with the `-fpen` option to control the handling of exceptions.

To determine the cause of the unaligned data when using idb, follow these steps:

1. Run the debugger, specifying the program with the unaligned data (shown as `testprog` in the following example): **idb testprog**
2. Before you run the program, enter the `catch unaligned` command:
   (idb) **catch unaligned**
3. Run the program:
   (idb) **run**
   Unaligned access pid=28413 <testprog> va=140000154
   pc=3ff80805d60
   ra=1200017e8 type=stl
   Thread received signal BUS
   stopped at [oops:13 0x120001834]
   13          end
4. Enter a `list` command to display the source code at line 12:
   (idb) **list 12**
   12          i4 = 1
   >  13          end
5. Enter the `where` command to find the location of the unaligned access:
   (idb) **where**
6. Use any other appropriate debugger commands needed to isolate the cause of the unaligned data, such as `up`, `list`, and `down`.
7. Repeat these steps for other areas where unaligned data is reported. Use the `rerun` command to run the program again instead of exiting the debugger and running it from the shell prompt.

8. After fixing the causes of the unaligned data, compile and link the program again.

# Data and I/O

**Data Representation**

# Data Representation Overview

See these topics:

Intrinsic Data Types

Integer Data Representations Overview

Logical Data Representations

Native IEEE Floating-Point Representations Overview

Character Representation

Hollerith Representation

# Intrinsic Data Types

Intel® Fortran expects numeric data to be in native little endian order, in which the least-significant, right-most zero bit (bit 0) or byte has a lower address than the most-significant, left-most bit (or byte). For information on using nonnative big endian and VAX* floating-point formats, see Converting Unformatted Numeric Data.

The symbol :A in any figure specifies the address of the byte containing bit 0, which is the starting address of the represented data element.

The following table lists the intrinsic data types used by Intel Fortran, the storage required, and valid ranges. For example, the declaration INTEGER(4) is the same as INTEGER(KIND=4) and INTEGER*4.

| Data Type | Storage | Description |
|---|---|---|
| BYTE<br>INTEGER(1) | 1 byte<br>(8 bits) | A BYTE declaration is a signed integer data type equivalent to INTEGER(1). |

| INTEGER | See INTEGER(2), INTEGER(4), and INTEGER(8). | Signed integer, either INTEGER(2), INTEGER(4), or INTEGER(8). The size is controlled by the -integer_size *nn* compiler option. The default is -integer_size 32 (INTEGER(4)). |
|---|---|---|
| INTEGER(1) | 1 byte (8 bits) | Signed integer value from -128 to 127. |
| INTEGER(2) | 2 bytes (16 bits) | Signed integer value from -32,768 to 32,767. |
| INTEGER(4) | 4 bytes (32 bits) | Signed integer value from -2,147,483,648 to 2,147,483,647. |
| INTEGER(8) | 8 bytes (64 bits) | Signed integer value from -9,223,372,036,854,775,808 to 9,223,372,036,854,775,807. |
| REAL(4) REAL | 4 bytes (32 bits) | Single-precision real floating-point values in IEEE S_floating format ranging from 1.17549435E-38 to 3.40282347E38. Values between 1.17549429E-38 and 1.40129846E-45 are denormalized (subnormal). |
| REAL(8) DOUBLE PRECISION | 8 bytes (64 bits) | Double-precision real floating-point values in IEEE T_floating format ranging from 2.2250738585072013D-308 to 1.7976931348623158D308. Values between 2.2250738585072008D-308 and 4.94065645841246544D-324 are denormalized (subnormal). |
| REAL(16) EXTENDED PRECISION | 16 bytes (128 bits) | Extended-precision real floating-point values in IEEE-style X_floating format ranging from 6.47517511943802511092443895822764655249996Q-4966 to 1.18973149535723176508575932662800701619647Q4932. |
| COMPLEX(4) COMPLEX | 8 bytes (64 bits) | Single-precision complex floating-point values in a pair of IEEE S_floating format parts: real and imaginary. The real and imaginary parts each range from 1.17549435E-38 to 3.40282347E38. Values between 1.17549429E-38 and 1.40129846E-45 are denormalized (subnormal). |
| COMPLEX(8) DOUBLE COMPLEX | 16 bytes (128 bits) | Double-precision complex floating-point values in a pair of IEEE T_floating format parts: real and imaginary. The real and imaginary parts each range from 2.2250738585072013D-308 to 1.7976931348623158D308. Values between 2.2250738585072008D-308 and 4.94065645841246544D-324 are denormalized (subnormal). |
| COMPLEX(16) EXTENDED PRECISION | 32 bytes (256 bits) | Extended-precision complex floating-point values in a pair of IEEE-style X_floating format parts: real and imaginary. The real and imaginary parts each range from 6.47517511943802511092443895822764655249996Q-4966 to 1.18973149535723176508575932662800701619647Q4932. |
| LOGICAL | See LOGICAL(2), LOGICAL(4), and LOGICAL(8). | Logical value, either LOGICAL(2), LOGICAL(4), or LOGICAL(8). The size is controlled by the -integer_size *nn* compiler option. The default is -integer_size 32 (LOGICAL(4)). |

| LOGICAL(1) | 1 byte (8 bits) | Logical values .TRUE. or .FALSE. |
|---|---|---|
| LOGICAL(2) | 2 bytes (16 bits) | Logical values .TRUE. or .FALSE. |
| LOGICAL(4) | 4 bytes (32 bits) | Logical values .TRUE. or .FALSE. |
| LOGICAL(8) | 8 bytes (64 bits) | Logical values .TRUE. or .FALSE. |
| CHARACTER | 1 byte (8 bits) per character | Character data represented by character code convention. Character declarations can be in the form CHARACTER(LEN=$n$) or CHARACTER*$n$, where $n$ is the number of bytes or $n$ is (*) to indicate passed-length format. |
| HOLLERITH | 1 byte (8 bits) per Hollerith character | Hollerith constants. |

In addition, you can define binary (bit) constants as explained in the *Language Reference.*

The following sections discuss the intrinsic data types in more detail:

- Integer Data Representations
- Logical Data Representations
- Native IEEE Floating-Point Representations
- Character Representation
- Hollerith Representation

**Integer Data Representations**

# Integer Data Representations Overview

Integer data lengths can be 1, 2, 4, or 8 bytes in length.

The default data size used for an INTEGER data declaration is INTEGER(4) (same as INTEGER(KIND=4)), unless the `-integer_size 16` or the `-integer_size 64` option was specified.

Integer data is signed with the sign bit being 0 (zero) for positive numbers and 1 for negative numbers.

The following sections discuss integer data:

- INTEGER(KIND=1) Representation
- INTEGER(KIND=2) Representation
- INTEGER(KIND=4) Representation
- INTEGER(KIND=8) Representation

# INTEGER(KIND=1) Representation

INTEGER(1) values range from -128 to 127 and are stored in 1 byte, as shown below:



ZK-9814-GE
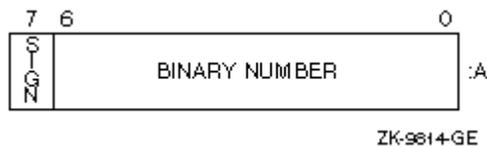
Integers are stored in a two's complement representation. For example:

```
+22 = 16(hex)
 -7 = F9(hex)
```

# INTEGER(KIND=2) Representation

INTEGER(2) values range from -32,768 to 32,767 and are stored in 2 contiguous bytes, as shown below:



ZK-0798-GE
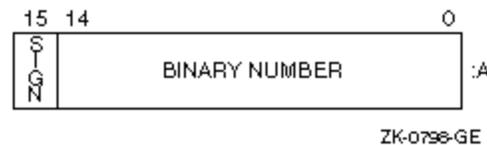
Integers are stored in a two's complement representation. For example:

```
+22 = 0016(hex)
 -7 = FFF9(hex)
```

# INTEGER(KIND=4) Representation

INTEGER(4) values range from -2,147,483,648 to 2,147,483,647 and are stored in 4 contiguous bytes, as shown below.
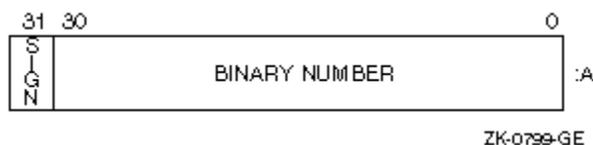


ZK-0799-GE

Integers are stored in a two's complement representation.

# INTEGER(KIND=8) Representation

INTEGER(8) values range from -9,223,372,036,854,775,808 to 9,223,372,036,854,775,807 and are stored in 8 contiguous bytes, as shown below.



ZK-5299A-GE
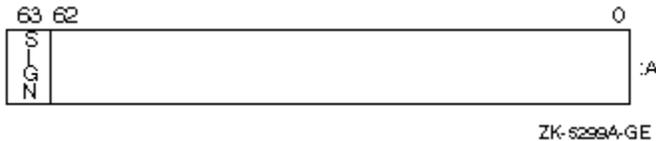
Integers are stored in a two's complement representation.

# Logical Data Representations

Logical data lengths can be 1-, 2-, 4-, or 8-bytes in length.

The default data size used for a LOGICAL data declaration is LOGICAL(4) (same as LOGICAL(KIND=4)), unless the `-integer_size 16` or `-integer_size 64` option was specified.

To improve performance on Itanium®-based systems, use LOGICAL(4) (or LOGICAL(8)) rather than LOGICAL(2) or LOGICAL(1).

LOGICAL(KIND=1) values are stored in 1 byte. In addition to having logical values .TRUE. and .FALSE., LOGICAL(1) data can also have values in the range -128 to 127. Logical variables can also be interpreted as integer data.

In addition to LOGICAL(1), logical values can also be stored in 2 (LOGICAL(2)), 4 (LOGICAL(4)), or 8 (LOGICAL(8)) contiguous bytes, starting on an arbitrary byte boundary.

If the `-fpscomp nological` compiler option is set (the default), the low-order bit determines whether the logical value is true or false. Specify `-fpscomp logical` for Microsoft* Fortran PowerStation logical values, where 0 (zero) is false and non-zero values are true.

LOGICAL(1), LOGICAL(2), LOGICAL(4), and LOGICAL(8) data representation (when `-fpscomp nological` option was set) appears below:

112
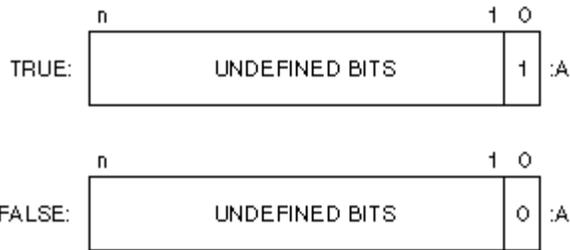
Key: n = 7, 15, 31, or 63 depending on LOGICAL declaration size

ZK-S300A-GE

**Native IEEE* Floating-Point Representations**

# Native IEEE* Floating-Point Representations Overview

The REAL(4) (IEEE* S_floating), REAL(8) (IEEE T_floating), and REAL(16) (IEEE-style X_floating) formats are stored in standard little endian IEEE binary floating-point notation. (See IEEE Standard 754 for additional information about IEEE binary floating point notation.) COMPLEX() formats use a pair of REAL values to denote the real and imaginary parts of the data.

All floating-point formats represent fractions in sign-magnitude notation, with the binary radix point to the right of the most-significant bit. Fractions are assumed to be normalized, and therefore the most-significant bit is not stored (this is called "hidden bit normalization"). This bit is assumed to be 1 unless the exponent is 0. If the exponent equals 0, then the value represented is denormalized (subnormal) or plus or minus zero.

Intrinsic REAL kinds are 4 (single precision), 8 (double precision), and 16 (extended precision), such as REAL(KIND=4) for single-precision floating-point data. Intrinsic COMPLEX kinds are also 4 (single precision), 8 (double precision), and 16 (extended precision).

To obtain the kind of a variable, use the KIND intrinsic function. You can also use a size specifier, such as REAL*4, but be aware this is an extension to the Fortran 95 standard.

If you omit certain compiler options, the default sizes for REAL and COMPLEX data declarations are as follows:

- For REAL data declarations without a kind parameter (or size specifier), the default size is REAL (KIND=4) (same as REAL*4).

- For COMPLEX data declarations without a kind parameter (or size specifier), the default data size is COMPLEX (KIND=4) (same as COMPLEX*8).

To control the size of all REAL or COMPLEX declarations without a kind parameter, use the `-real_size 64` or `-real_size 128` options; the default is `-real_size 32`.

You can explicitly declare the length of a REAL or a COMPLEX declaration using a kind parameter, or specify DOUBLE PRECISION or DOUBLE COMPLEX. To control the size of all DOUBLE PRECISION and DOUBLE COMPLEX declarations, use the `-double_size 128` option; the default is `-double_size 64`.

The following sections discuss floating-point data:

- REAL(KIND=4) (Single Precision)  Representation
- REAL(KIND=8) (Double Precision) Representation
- REAL(KIND=16) (Extended Precision) Representation
- COMPLEX(KIND=4) (Single Precision) Representation
- COMPLEX(KIND=8) (Double Precision) Representation
- COMPLEX(KIND=16 Representation

For information on reading or writing floating-point data other than native IEEE little endian data, see Converting Unformatted Numeric Data.

See also File fordef.for and Its Usage.

# REAL(KIND=4) (REAL) Representation

REAL(4) (same as REAL(KIND=4)) data occupies 4 contiguous bytes stored in IEEE S_floating format. Bits are labeled from the right, 0 through 31, as shown below.



The form of REAL(4) data is sign magnitude, with bit 31 the sign bit (0 for positive numbers, 1 for negative numbers), bits 30:23 a binary exponent in excess 127 notation, and bits 22:0 a normalized 24-bit fraction including the redundant most-significant fraction bit not represented.

The value of data is in the approximate range: 1.17549435E-38 (normalized) to 3.40282347E38. The IEEE denormalized (subnormal) limit is 1.40129846E-45. The precision is approximately one part in 2**23; typically 7 decimal digits.

# REAL(KIND=8) (DOUBLE PRECISION) Representation

REAL(8) (same as REAL(KIND=8)) data occupies 8 contiguous bytes stored in IEEE T_floating format. Bits are labeled from the right, 0 through 63, as shown below.



ZK-9816-GE
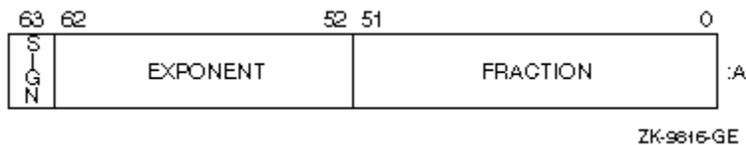
The form of REAL(8) data is sign magnitude, with bit 63 the sign bit (0 for positive numbers, 1 for negative numbers), bits 62:52 a binary exponent in excess 1023 notation, and bits 51:0 a normalized 53-bit fraction including the redundant most-significant fraction bit not represented.

The value of data is in the approximate range: 2.2250738585072013D-308 (normalized) to 1.7976931348623158D308. The IEEE denormalized (subnormal) limit is 4.94065645841246544D-324. The precision is approximately one part in 2**52; typically 15 decimal digits.

# REAL(KIND=16) (EXTENDED PRECISION) Representation

REAL(16) (same as REAL(KIND=16)) data occupies 16 contiguous bytes stored in IEEE-style X_floating format. Bits are labeled from the right, 0 through 127, as shown below.



ZK-7420A-GE
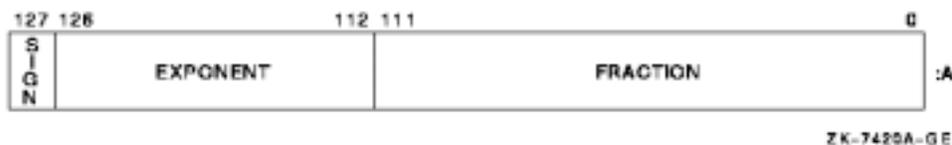
The form of REAL(16) data is sign magnitude, with bit 127 the sign bit (0 for positive numbers, 1 for negative numbers), bits 126:112 a binary exponent in

excess 16383 notation, and bits 111:0 a normalized 113-bit fraction including the redundant most-significant fraction bit not represented.

The value of data is in the approximate range: 6.4751751194380251109244389582276465524996Q-4966 to 1.1897314953572317650857593266280070161964770Q4932. Unlike other floating-point formats, there is little if any performance penalty from using denormalized extended-precision numbers. This is because accessing denormalized REAL (KIND=16) numbers does not result in an arithmetic trap (the extended-precision format is emulated in software). The smallest normalized number is 3.3621031431120935062626778173217531753Q-4932.

The precision is approximately one part in 2**112 or typically 33 decimal digits.

# COMPLEX(KIND=4) (COMPLEX) Representation

COMPLEX(4) (same as COMPLEX(KIND=4) and COMPLEX*8) data is 8 contiguous bytes containing a pair of REAL(4) values stored in IEEE S_floating format. The low-order 4 bytes contain REAL(4) data that represents the real part of the complex number. The high-order 4 bytes contain REAL(4) data that represents the imaginary part of the complex number, as shown below.
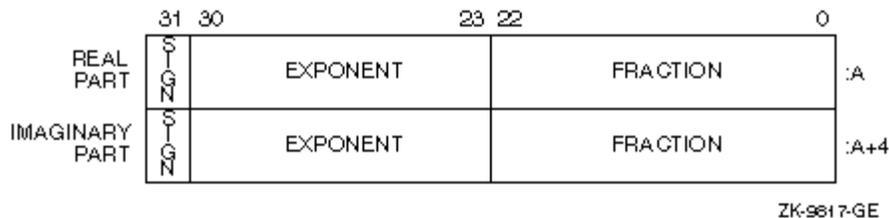


The limits and underflow characteristics for REAL(4) apply to the two separate real and imaginary parts of a COMPLEX(4) number. Like REAL(4) numbers, the sign bit representation is 0 (zero) for positive numbers and 1 for negative numbers.

# COMPLEX(KIND=8) (DOUBLE COMPLEX) Representation

COMPLEX(8) (same as COMPLEX(KIND=8) and COMPLEX*16) data is 16 contiguous bytes containing a pair of REAL(8) values stored in IEEE T_floating format. The low-order 8 bytes contain REAL(8) data that represents the real part

of the complex data. The high-order 8 bytes contain REAL(8) data that represents the imaginary part of the complex data, as shown below.



ZK-9818-GE
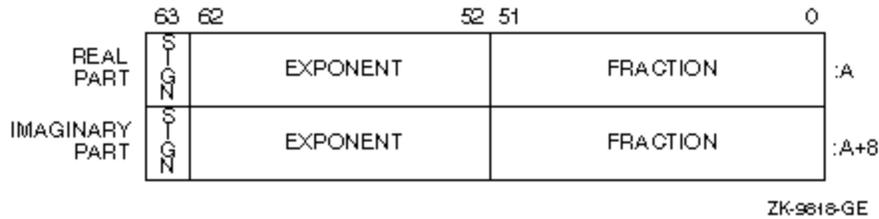
The limits and underflow characteristics for REAL(8) apply to the two separate real and imaginary parts of a COMPLEX(8) number. Like REAL(8) numbers, the sign bit representation is 0 (zero) for positive numbers and 1 for negative numbers.

# COMPLEX(KIND=16) Representation

COMPLEX(16) (same as COMPLEX(KIND=16) or COMPLEX*32) data is 32 contiguous bytes containing a pair of REAL(16) values stored in IEEE-style X_floating format. The low-order 16 bytes contain REAL(16) data that represents the real part of the complex data. The high-order 16 bytes contain REAL(8) data that represents the imaginary part of the complex data, as shown below.



LJ–06690
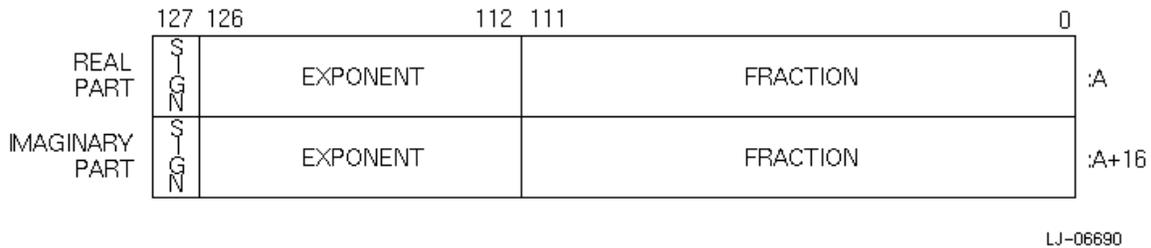
The limits and underflow characteristics for REAL(16) apply to the two separate real and imaginary parts of a COMPLEX(16) number. Like REAL(16) numbers, the sign bit representation is 0 (zero) for positive numbers and 1 for negative numbers.

# File fordef.for and Its Usage

The parameter file `/opt/intel_fc_80/include/fordef.for` contains symbols and INTEGER*4 values corresponding to the classes of floating-point representations. Some of these classes are exceptional ones such as bit patterns that represent positive denormalized numbers.

With this file of symbols and with the FP_CLASS intrinsic function, you have the flexibility of identifying exceptional numbers so that, for example, you can replace positive and negative denormalized numbers with true zero.

The following is a simple example of identifying floating-point bit representations:

```fortran
include 'fordef.for'
real*4 a
integer*4 class_of_bits
a = 57.0    ! Bit pattern is a finite number
class_of_bits = fp_class(a)
if ( class_of_bits .eq. for_k_fp_pos_norm  .or.   &
     class_of_bits .eq. for_k_fp_neg_norm       ) then
   print *, a, ' is a non-zero and non-exceptional value'
else
   print *, a, ' is zero or an exceptional value'
end if
end
```

In this example, the symbol `for_k_fp_pos_norm` in the file `/opt/intel_fc_80/include/fordef.for` plus the REAL*4 value 57.0 to the FP_CLASS intrinsic function results in the execution of the first print statement.

The table below explains the symbols in the file `/opt/intel_fc_80/include/fordef.for` and their corresponding floating-point representations.

**Symbols in File fordef.for**

| Symbol Name | Class of Floating-Point Bit Representation |
|---|---|
| FOR_K_FP_SNAN | Signaling NaN |
| FOR_K_FP_QNAN | Quiet NaN |
| FOR_K_FP_POS_INF | Positive infinity |
| FOR_K_FP_NEG_INF | Negative infinity |
| FOR_K_FP_POS_NORM | Positive normalized finite number |
| FOR_K_FP_NEG_NORM | Negative normalized finite number |
| FOR_K_FP_POS_DENORM | Positive denormalized number |
| FOR_K_FP_NEG_DENORM | Negative denormalized number |
| FOR_K_FP_POS_ZERO | Positive zero |
| FOR_K_FP_NEG_ZERO | Negative zero |

Another example of using file `fordef.for` and intrinsic function FP_CLASS follows. The goals of this program are to quickly read any 32-bit pattern into a REAL*4 number from an unformatted file with no exception reporting and to replace denormalized numbers with true zero:

```
include 'fordef.for'
      real*4 a(100)
      integer*4 class_of_bits
!     open an unformatted file as unit 1
!     ...
      read (1) a
      do i = 1, 100
         class_of_bits = fp_class(a(i))
         if ( class_of_bits .eq. for_k_fp_pos_denorm  .or.
  &
            class_of_bits .eq. for_k_fp_neg_denorm       )
then
            a(i) = 0.0
         end if
      end do
      close (1)
      end
```

You can compile this program with any value of `-fpen`. Intrinsic function
FP_CLASS helps to find and replace denormalized numbers with zeroes before
the program can attempt to perform calculations on the denormalized numbers.
On the other hand, if this program did not replace denormalized numbers read
from unit 1 with zeroes and the program was compiled with `-fpe0`, then the first
attempted calculation on a denormalized number would result in a floating-point
exception.

File `fordef.for` and intrinsic function FP_CLASS can work together to identify
NaNs. A variation of the previous example would contain the symbols
`for_k_fp_snan` and `for_k_fp_qnan` in the IF statement. A faster way to do this is
based on the intrinsic function ISNAN. One modification of the previous example,
using ISNAN, follows:

```
!     The ISNAN function does not need file fordef.for
      real*4 a(100)
!     open an unformatted file as unit 1
!     ...
      read (1) a
      do i = 1, 100
         if ( isnan (a(i)) ) then
            print *, 'Element ', i, ' contains a NaN'
         end if
      end do
      close (1)
      end
```

You can compile this program with any value of `-fpen`.

# Character Representation

A character string is a contiguous sequence of bytes in memory, as shown below.

```
┌──────────────┐
│   CHAR 1     │    :A
└──────────────┘
        •
        •
        •
┌──────────────┐
│   CHAR L     │    :A+L-1
└──────────────┘
     ZK-0809-GE
```
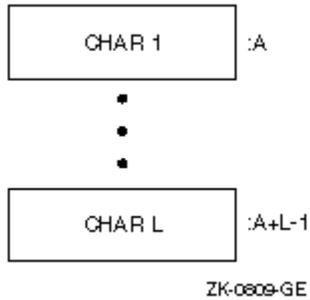
A character string is specified by two attributes: the address A of the first byte of the string, and the length L of the string in bytes. The length L of a string is in the range 1 through 2,147,483,647 $(2**31-1)$ .

# Hollerith Representation

Hollerith constants are stored internally, one character per byte, as shown below.
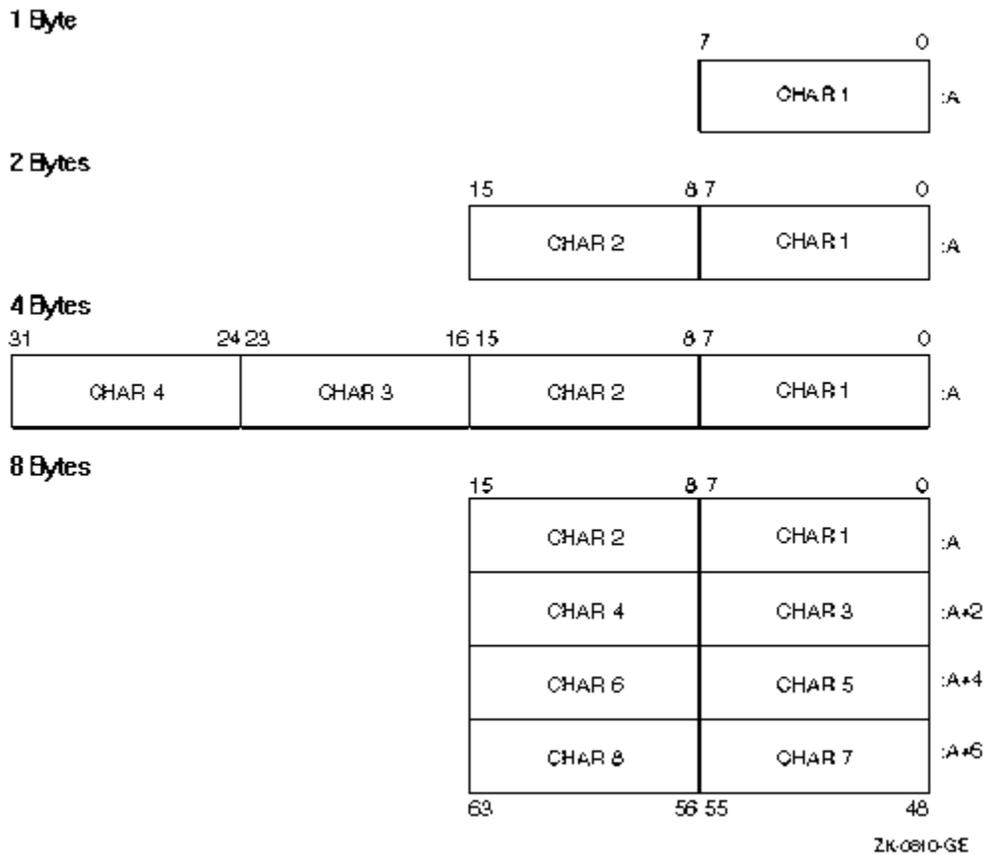
**Converting Unformatted Data**

# Converting Unformatted Data Overview

This section describes how you can use Intel® Fortran to read and write nonnative unformatted numeric data.

See these topics:

Supported Native and Nonnative Numeric Formats

Limitations of Numeric Conversion

Methods of Specifying the Data Format: Overview

Porting Nonnative Data

# Supported Native and Nonnative Numeric Formats

Intel® Fortran supports the following little endian floating-point formats in memory:

| Floating-point size | Format in memory |
|---|---|
| REAL(KIND=4) COMPLEX(KIND=4) | IEEE* S_floating |
| REAL(KIND=8) COMPLEX(KIND=8) | IEEE T_floating |
| REAL(KIND=16 COMPLEX(KIND=16) | IEEE X_floating |

If your program needs to read or write unformatted data files containing a floating-point format that differs from the format in memory for that data size, you can request that the unformatted data be converted.

Data storage in different computers uses a convention of either little endian or big endian storage. The storage convention generally applies to numeric values that span multiple bytes, as follows:

- Little endian storage occurs when:
    - The least significant bit (LSB) value is in the byte with the lowest address.
    - The most significant bit (MSB) value is in the byte with the highest address.
    - The address of the numeric value is the byte containing the LSB. Subsequent bytes with higher addresses contain more significant bits.
- Big endian storage occurs when:
    - The least significant bit (LSB) value is in the byte with the highest address.
    - The most significant bit (MSB) value is in the byte with the lowest address.
    - The address of the numeric value is the byte containing the MSB. Subsequent bytes with higher addresses contain less significant bits.

The following figure shows the difference between the two byte-ordering schemes:

**Little and Big Endian Storage of an INTEGER Value**

ZK-6654A-GE
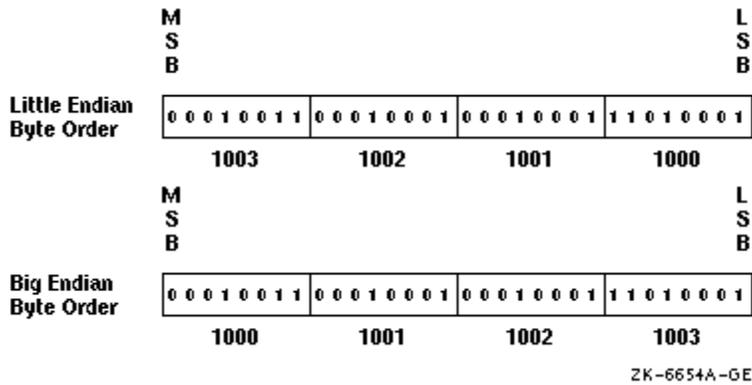
Moving unformatted data files between big endian and little endian computers requires that the data be converted.

Intel Fortran provides the capability for programs to read and write unformatted data (originally written using unformatted I/O statements) in several nonnative floating-point formats and in big endian INTEGER or floating-point format. Supported nonnative floating-point formats include VAX* little endian floating-point formats supported by VAX FORTRAN, standard IEEE big endian floating-point format found on most Sun Microsystems* systems and IBM RISC* System/6000 systems, IBM floating-point formats (associated with the IBM's System/370 and similar systems), and CRAY* floating-point formats.

Converting unformatted data instead of formatted data is generally faster and is less likely to lose precision of floating-point numbers.

The native memory format includes little endian integers and little endian IEEE floating-point formats, S_floating for REAL(KIND=4) and COMPLEX(KIND=4) declarations, T_floating for REAL(KIND=8) and COMPLEX(KIND=8) declarations, and IEEE X_floating for REAL(KIND=16) and COMPLEX(KIND=16) declarations.

The keywords for supported nonnative unformatted file formats and their data types are listed in the following table:

**Nonnative Numeric Format Keywords and Supported Data Types**

| Keyword | Description |
|---|---|
| BIG_ENDIAN | Big endian integer data of the appropriate INTEGER size (one, two, four, or eight bytes) and big endian IEEE floating-point formats for REAL and COMPLEX single- and double- and extended-precision numbers. INTEGER (KIND=1) or INTEGER*1 data is the same for little endian and big endian. |
| CRAY | Big endian integer data of the appropriate INTEGER size |

| | (one, two, four, or eight bytes) and big endian CRAY proprietary floating-point format for REAL and COMPLEX single- and double-precision numbers. |
|---|---|
| FDX | Native little endian integers of the appropriate INTEGER size (one, two, four, or eight bytes) and the following little endian proprietary floating-point formats:<br><br>• VAX F_float for REAL (KIND=4) and COMPLEX (KIND=4)<br>• VAX D_float for REAL (KIND=8) and COMPLEX (KIND=8)<br>• IEEE style X_float for REAL (KIND=16) and COMPLEX (KIND=16) |
| FGX | Native little endian integers of the appropriate INTEGER size (one, two, four, or eight bytes) and the following little endian proprietary floating-point formats:<br><br>• VAX F_float for REAL (KIND=4) and COMPLEX (KIND=4)<br>• VAX G_float for REAL (KIND=8) and COMPLEX (KIND=8)<br>• IEEE style X_float for REAL (KIND=16) and COMPLEX (KIND=16) |
| IBM | Big endian integer data of the appropriate INTEGER size (one, two, or four bytes) and big endian IBM proprietary (System\370 and similar) floating-point format for REAL and COMPLEX single- and double-precision numbers. |
| LITTLE_ENDIAN | Native little endian integers of the appropriate INTEGER size (one, two, four, or eight bytes) and the following native little endian IEEE floating-point formats:<br><br>• S_float for REAL (KIND=4) and COMPLEX (KIND=4)<br>• T_float for REAL (KIND=8) and COMPLEX (KIND=8)<br>• IEEE style X_float for REAL (KIND=16) and COMPLEX (KIND=16) |
| NATIVE | No conversion occurs between memory and disk. This is the default for unformatted files. |
| VAXD | Native little endian integers of the appropriate INTEGER size (one, two, four, or eight bytes) and the following little |

| | |
|---|---|
| | endian VAX proprietary floating-point formats:<br><br>• VAX F_float for REAL (KIND=4) and COMPLEX (KIND=4)<br>• VAX D_float for REAL (KIND=8) and COMPLEX (KIND=8)<br>• VAX H_float for REAL (KIND=16) and COMPLEX (KIND=16) |
| VAXG | Native little endian integers of the appropriate INTEGER size (one, two, four, or eight bytes) and the following little endian VAX proprietary floating-point formats:<br><br>• VAX F_float for REAL (KIND=4) and COMPLEX (KIND=4)<br>• VAX G_float for REAL (KIND=8) and COMPLEX (KIND=8)<br>• VAX H_float for REAL (KIND=16) and COMPLEX (KIND=16) |

When reading a nonnative format, the nonnative format on disk is converted to native format in memory. If a converted nonnative value is outside the range of the native data type, a run-time message is displayed.

# Limitations of Numeric Conversion

The Intel Fortran floating-point conversion solution is not expected to fulfill all floating-point conversion needs.

For instance, data fields in record structure variables (specified in a STRUCTURE statement) and data components of derived types (TYPE statement) are not converted. When they are later examined as separate fields by the program, they will remain in the binary format they were stored in on disk, unless the program is modified. With EQUIVALENCE statements, the data type of the variable named in the I/O statement is used.

If a program reads an I/O record containing multiple format floating-point fields into a single variable (such as an array) instead of their respective variables, the fields will not be converted. When they are later examined as separate fields by the program, they will remain in the binary format they were stored in on disk, unless the program is modified.

**Methods of Specifying the Data Format**

# Methods of Specifying the Data Format: Overview

There are six methods of specifying a nonnative numeric format for unformatted data:

- Setting an environment variable for a specific unit number before the file is opened. The environment variable is named FORT_CONVERT*n*, where *n* is the unit number. See Environment Variable FORT_CONVERTn Method.
- Setting an environment variable for a specific file name extension before the file is opened. The environment variable is named FORT_CONVERT.*ext* or FORT_CONVERT_*ext*, where *ext* is the file name extension (suffix). See Environment Variable FORT_CONVERT.ext or FORT_CONVERT_ext Method.
- Setting an environment variable for a set of units before any files are opened. The environment variable is named F_UFMTENDIAN. See Environment Variable F_UFMTENDIAN Method.
- Specifying the CONVERT keyword in the OPEN statement for a specific unit number. See OPEN Statement CONVERT Method.
- Compiling the program with an OPTIONS statement that specifies the CONVERT=*keyword* qualifier. This method affects all unit numbers using unformatted data specified by the program. See OPTIONS Statement Method.
- Compiling the program with the command-line `-convert keyword` option, which affects all unit numbers that use unformatted data specified by the program. See Compiler Option -convert Method.

If none of these methods are specified, the native LITTLE_ENDIAN format is assumed (no conversion occurs between disk and memory).

Any keyword listed in Supported Native and Nonnative Numeric Formats can be used with any of these methods, except for the Environment Variable F_UFMTENDIAN Method, which supports only LITTLE_ENDIAN and BIG_ENDIAN.

If you specify more than one method, the order of precedence when you open a file with unformatted data is to:

1. Check for an environment variable (FORT_CONVERT*n*) for the specified unit number (applies to any file opened on a particular unit).
2. Check for an environment variable (FORT_CONVERT.*ext* is checked before FORT_CONVERT_*ext*) for the specified file name extension (applies to all files opened with the specified file name extension).
3. Check for an environment variable (F_UFMTENDIAN) for the specified unit number (or for all units).
4. Check the OPEN statement CONVERT qualifier.

5.  Check whether an OPTIONS statement with a CONVERT=*keyword* qualifier was present when the program was compiled.
6.  Check whether the compiler option `-convert keyword` was present when the program was compiled.

# Environment Variable FORT_CONVERTn Method

You can use this method to specify a non-native numeric format for each specified unit number. You specify the numeric format at run time by setting the appropriate environment variable before an implicit or explicit OPEN to that unit number.

When the appropriate environment variable is set when you open the file, the environment variable is always used because this method takes precedence over the other methods. For instance, you might use this method to specify that a unit number will use a particular format instead of the format specified in the program (perhaps for a one-time file conversion).

For example, assume you have a previously compiled program that reads numeric data from unit 28 and writes it to unit 29 using unformatted I/O statements. You want the program to read nonnative big endian (IEEE floating-point) format from unit 28 and write that data in native little endian format to unit 29. In this case, the data is converted from big endian IEEE format to native little endian IEEE memory format when read from unit 28, and then written without conversion in native little endian IEEE format to unit 29.

Without requiring source code modification or recompilation of this program, the following command sequence sets the appropriate environment variables before running the program (`/usr/users/leslie/convieee`):

```
setenv FORT_CONVERT28 BIG_ENDIAN
setenv FORT_CONVERT29 NATIVE
/usr/users/leslie/convieee
```

The following figure shows the data formats used on disk and in memory when the example file (`/usr/users/leslie/convieee`) is run after the environment variables are set.

**Sample Unformatted File Conversion**
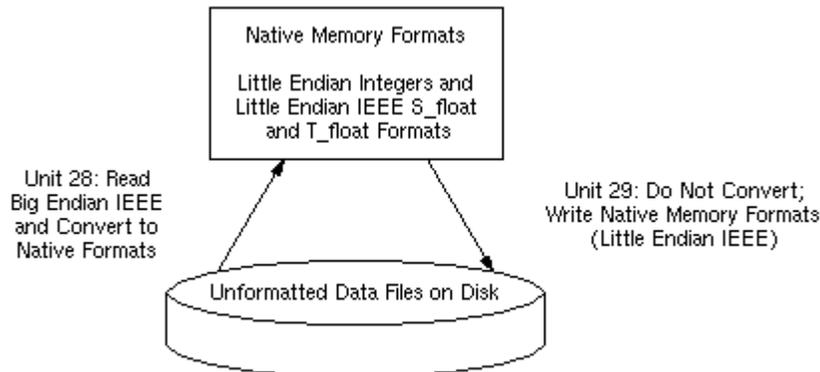
ZK-8326A-GE

This method takes precedence over other methods.

# Environment Variable FORT_CONVERT.ext or FORT_CONVERT_ext Method

You can use this method to specify a non-native numeric format for each specified file name extension (suffix). You specify the numeric format at run time by setting the appropriate environment variable before an implicit or explicit OPEN to one or more unformatted files. You can use the format FORT_CONVERT.*ext* or FORT_CONVERT_*ext* (where *ext* is the file extension or suffix). The FORT_CONVERT.*ext* environment variable is checked before FORT_CONVERT_*ext* environment variable (if *ext* is the same).

For example, assume you have a previously compiled program that reads numeric data from one file and writes to another file using unformatted I/O statements. You want the program to read nonnative big endian (IEEE floating-point) format from a file with a .dat file extension extension and write that data in native little endian format to a file with a extension of .data. In this case, the data is converted from big endian IEEE format to native little endian IEEE memory format (S_float and T_float) when read from file.dat, and then written without conversion in native little endian IEEE format to the file with a suffix of .data, assuming that environment variables FORT_CONVERT.DATA and FORT_CONVERT*n* (for that unit number) are not defined.

Without requiring source code modification or recompilation of this program, the following command sequence sets the appropriate environment variables before running the program:

```
setenv FORT_CONVERT.DAT BIG_ENDIAN
/usr/users/proj2/cvbigend
```

The FORT_CONVERT*n* method takes precedence over this method. When the appropriate environment variable is set when you open the file, the FORT_CONVERT.*ext* or FORT_CONVERT_*ext* environment variable is used if a FORT_CONVERT*n* environment variable is not set for the unit number.

The FORT_CONVERT*n* and the FORT_CONVERT.*ext* or FORT_CONVERT_*ext* environment variable methods take precedence over the other methods. For instance, you might use this method to specify that a unit number will use a particular format instead of the format specified in the program (perhaps for a one-time file conversion).

You can set the appropriate environment variable using the format FORT_CONVERT.*ext* or FORT_CONVERT_*ext*. Consider using the FORT_CONVERT_*ext* form, because a dot (.) cannot be used for environment variable names on certain Linux* command shells. If you do define both FORT_CONVERT.*ext* and FORT_CONVERT_*ext* for the same extension (*ext*), the file defined by FORT_CONVERT.*ext* is used.

# Environment Variable F_UFMTENDIAN Method

This little-endian-to-big-endian conversion feature is intended for Fortran unformatted input/output operations. It enables the development and processing of files with big-endian data organization.

The feature also enables processing of the files developed on processors that accept big-endian data format and producing the files for such processors on IA-32-based little-endian systems.

This little-endian-to-big-endian conversion is accomplished by the following operations:

- The WRITE operation converts little endian format to big endian format.
- The READ operation converts big endian format to little endian format.

## Little-to-Big Endian Conversion Environment Variable

In order to use the little-endian-to-big-endian conversion feature, specify the numbers of the units to be used for conversion purposes by setting the F_UFMTENDIAN environment variable. Then, the READ/WRITE statements that

use these unit numbers, will perform relevant conversions. Other READ/WRITE statements will work in the usual way.

In the general case, the variable consists of two parts divided by a semicolon. No spaces are allowed inside the F_UFMTENDIAN value. The variable has the following syntax:

```
F_UFMTENDIAN=MODE | [MODE;] EXCEPTION
```

where:

```
MODE = big | little
EXCEPTION = big:ULIST | little:ULIST | ULIST
ULIST = U | ULIST,U
U = decimal | decimal -decimal
```

- MODE defines current format of data, represented in the files; it can be omitted.
  The keyword little means that the data has little endian format and will not be converted. This is the default.
  The keyword big means that the data has big endian format and will be converted.
- EXCEPTION is intended to define the list of exclusions for MODE.
  EXCEPTION keyword (little or big) defines data format in the files that are connected to the units from the EXCEPTION list. This value overrides MODE value for the units listed. The EXCEPTION keyword and the colon can be omitted. The default when the keyword is omitted is big.
- Each list member U is a simple unit number or a number of units. The number of list members is limited to 64.
  decimal is a non-negative decimal number less than $2^{32}$.

The command line for the variable setting in the shell is:

```
Sh: export F_UFMTENDIAN=MODE;EXCEPTION
```

## Note

The environment variable value should be enclosed in quotes if the semicolon is present.

## Another Possible Environment Variable Setting

The environment variable can also have the following syntax:

```
F_UFMTENDIAN=u[,u] . . .
```

The command line for the variable setting in the shell is:

Sh: **export  F_UFMTENDIAN=u[,u] . . .**

# Usage Examples

1. `F_UFMTENDIAN=big`

    All input/output operations perform conversion from big-endian to little-endian on `READ` and from little-endian to big-endian on `WRITE`.

2. `F_UFMTENDIAN="little;big:10,20"`
or `F_UFMTENDIAN=big:10,20`
or `F_UFMTENDIAN=10,20`

    In this case, only on unit numbers 10 and 20 the input/output operations perform big-little endian conversion.

3. `F_UFMTENDIAN="big;little:8"`

    In this case, on unit number 8 no conversion operation occurs. On all other units, the input/output operations perform big-little endian conversion.

4. `F_UFMTENDIAN=10-20`

    Define 10, 11, 12 … 19, 20 units for conversion purposes; on these units, the input/output operations perform big-little endian conversion.

5. Assume you set `F_UFMTENDIAN=10,100` and run the following program.

```
integer*4   cc4
integer*8   cc8
integer*4   c4
integer*8   c8
c4 = 456
c8 = 789

C  prepare little endian representation of data

open(11,file='lit.tmp',form='unformatted')
write(11) c8
write(11) c4
close(11)

C  prepare big endian representation of data

open(10,file='big.tmp',form='unformatted')
```

```
write(10) c8
write(10) c4
close(10)

C  read big endian data and operate with them on
C  little endian machine

open(100,file='big.tmp',form='unformatted')
read(100) cc8
read(100) cc4

C    Any operation with data, which have been read

C    . . .
close(100)
stop
end
```

Now compare `lit.tmp` and `big.tmp` files with the help of `od` utility.

```
> od -t x4 lit.tmp
0000000 00000008 00000315 00000000 00000008
0000020 00000004 000001c8 00000004
0000034

> od -t x4 big.tmp
0000000 08000000 00000000 15030000 08000000
0000020 04000000 c8010000 04000000
0000034
```

You can see that the byte order is different in these files.

# OPEN Statement CONVERT Method

You can use this method to specify a non-native numeric format for each specified unit number. This method requires an explicit file OPEN statement to specify the numeric format of the file for that unit number.

This method takes precedence over the OPTIONS statement and the compiler option `-convert keyword` method, but has a lower precedence than the environment variable methods.

For example, the following source code shows how the OPEN statement would be coded to read unformatted VAXD numeric data from unit 15, which might be processed and possibly written in native little endian format to unit 20. (Tthe absence of the CONVERT keyword or environment variables `FORT_CONVERT20`,

FORT_CONVERT.dat, FORT_CONVERT_dat, or F_UFMTENDIAN indicates native little endian data for unit 20.)

```
OPEN (CONVERT='VAXD', FILE='graph3.dat', FORM='UNFORMATTED', UNIT=15

.

OPEN (FILE='graph3_t.dat', FORM='UNFORMATTED', UNIT=20)
```

A hard-coded OPEN statement CONVERT *keyword* value cannot be changed after compile time. However, to allow selection of a particular format at run time, equate the CONVERT keyword to a variable and provide the user with a menu that allows selection of the appropriate format (menu choice sets the variable) before the OPEN occurs.

You can also select a particular format at run time for a unit number by using one of the environment variable methods (FORT_CONVERT*n*, FORT_CONVERT.*ext*, FORT_CONVERT_*ext*, or F_UFMTENDIAN), which take precedence over the OPEN statement CONVERT *keyword* method.

# OPTIONS Statement Method

You can only specify one numeric file format for all unformatted file unit numbers using this method unless you also use one of the environment variable methods or OPEN statement CONVERT *keyword* method.

You specify the numeric format at compile time and must compile all routines under the same OPTIONS statement CONVERT *keyword* qualifier. You could use one source program and compile it using different `ifort` commands to create multiple executable programs that each read a certain format.

The environment variable methods and the OPEN statement CONVERT method take precedence over this method. For instance, you might use the environment variable FORT_CONVERTn method or OPEN statement CONVERT method to specify each unit number that will use a format other than that specified using the `ifort` option method.

This method takes precedence over the `convert` *keyword* compiler option method.

You can use OPTIONS statements to specify the appropriate floating-point formats (in memory and in unformatted files) instead of using the corresponding ifort command qualifiers. For example, to use VAX F_floating and G_floating as the unformatted file format, specify the following OPTIONS statement:

```
OPTIONS /CONVERT=VAXG
```

Because this method affects all unit numbers, you cannot read data in one format and write it in another format, unless you use it in combination with one of the environment variable methods or the OPEN statement CONVERT keyword method to specify a different format for a particular unit number.

# Compiler Option -convert Method

You can only specify one numeric format for all unformatted file unit numbers using the compiler option `-convert` method unless you also use one (or more) of the previous methods.

You specify the numeric format at compile time and must compile all routines under the same `-convert keyword` compiler option. You could use the same source program and compile it using different `ifort` commands to create multiple executable programs that each read a certain format.

If you specify other methods, they take precedence over this method. For instance, you might use the environment variable or OPEN statement CONVERT keyword method to specify each unit number that will use a format different than that specified using the `-convert keyword` compiler option method for all other unit numbers.

For example, the following command compiles program `file.for` to use VAX D_floating (and F_floating) floating-point data for all unit numbers (unless superseded by one of the other methods). Data is converted between the file format and the little endian memory format (little endian integers, S_float and T_float little endian IEEE floating-point format). The created file, `vconvert.exe`, can then be run:

**ifort file.for -o vconvert -convert vaxd**

```
Because this method affects all unformatted file unit numbers, you
cannot read data in one format and write it in another file format
using the -convert keyword compiler option method alone. You can if you
use it in combination with the environment variable methods or the OPEN
statement CONVERT keyword method to specify a different format for a
particular unit number.
```

# Porting Nonnative Data

Keep this information in mind when porting nonnative data:

- When porting source code along with the unformatted data, vendors might use different units for specifying the record length (RECL specifier) of unformatted files. While formatted files are specified in units of characters (bytes), unformatted files are specified in longword units for Intel Fortran (default) and some other vendors.

  To allow you to specify the RECL units (bytes or longwords) for unformatted files without source file modification, use the `-assume byterecl` compiler option.

  The Fortran 90 standard (American National Standard Fortran 90, ANSI X3.198-1991, and International Standards Organization standard ISO/IEC 1539:1991), in Section 9.3.4.5, states: "If the file is being connected for unformatted input/output, the length is measured in processor-dependent units."

- Certain vendors apply different OPEN statement defaults to determine the record type. The default record type (RECORDTYPE) with Intel Fortran depends on the values for the ACCESS and FORM specifiers for the OPEN statement.
- Certain vendors use a different identifier for the logical data types, such as hex FF instead of 01 to denote "true."
- Source code being ported may be coded specifically for big endian use.

**Fortran I/O**

# Fortran I/O Overview

See these topics:

Logical I/O Units

Types of I/O Statements

Forms of I/O Statements

Files and File Characteristics Overview

Accessing and Assigning Files

Default Pathnames and File Names

Using Preconnected Standard I/O Files

Opening Files: OPEN Statement

# Logical I/O Units

In Intel Fortran, a *logical unit* is a channel through which data transfer occurs between the program and a device or file. You identify each logical unit with a logical unit number, which can be any nonnegative integer from 0 to a maximum value of 2,147,483,647 (2**31-1). For example:

```
READ (2,100) I,X,Y
```

This READ statement specifies that data is to be entered from the device or file corresponding to logical unit 2, in the format specified by the FORMAT statement labeled 100. When opening a file, use the UNIT specifier to indicate the unit number.

Fortran programs are inherently device-independent. The association between the logical unit number and the physical file can occur at run-time. Instead of changing the logical unit numbers specified in the source program, you can change this association at run time to match the needs of the program and the available resources. For example, before running the program, a script file can set the appropriate environment variable or allow the terminal user to type a directory path, file name, or both.

Use the same logical unit number specified in the OPEN statement for other I/O statements to be applied to the opened file, such as READ and WRITE.

The OPEN statement connects a unit number with an external file and allows you to explicitly specify file attributes and run-time options using OPEN statement specifiers (all files except internal files are called external files).

Certain unit numbers are preconnected to standard devices. Unit number 5 is associated with `stdin`, unit 6 with `stdout`, and unit 0 with `stderr`. At run time, if units 5 and 6 are specified by a record I/O statement (such as READ or WRITE) without having been explicitly opened by an OPEN statement, Intel Fortran implicitly opens units 5, 6, and 0 and associates them with their

respective operating system standard I/O files (if the corresponding FORT*n* environment variable is not set).

# Types of I/O Statements

The table below lists the Intel Fortran I/O statements:

| Category and statement name | Description |
|---|---|
| **File connection** | |
| OPEN | Connects a unit number with an external file and specifies file connection characteristics. |
| CLOSE | Disconnects a unit number from an external file. |
| **File inquiry** | |
| DEFINE FILE | Specifies file characteristics for a direct access relative file and connects the unit number to the file, similar to an OPEN statement. Provided for compatibility with compilers older than FORTRAN-77. |
| INQUIRE | Returns information about a named file, a connection to a unit, or the length of an output item list. |
| **Record position** | |
| BACKSPACE | Moves the record position to the beginning of the previous record (sequential access only). |
| DELETE | Marks a record at the current record position in a relative file as deleted (direct access only). |
| ENDFILE | Writes an end-of-file marker after the current record (sequential access only). |
| FIND | Changes the record position in a direct access file. Provided for compatibility with compilers older than FORTRAN-77. |
| REWIND | Sets the record position to the beginning of the file (sequential access only). |
| **Record input** | |
| READ | Transfers data from an external file record or an internal file to internal storage. |
| UNLOCK | Frees a record in a relative or sequential file that was locked by a previous READ statement. |
| ACCEPT | Reads input from stdin. Unlike READ, ACCEPT only provides formatted sequential input and does not specify a unit number. |
| **Record output** | |
| WRITE | Transfers data from internal storage to an external file record or to an internal file. |

| REWRITE | Transfers data from internal storage to an external file record at the current record position (direct access relative files only). |
|---|---|
| TYPE | Writes record output to `stdout` (same as PRINT). |
| PRINT | Transfers data from internal storage to `stdout`. Unlike WRITE, PRINT only provides formatted sequential output and does not specify a unit number. |

In addition to the READ, WRITE, REWRITE, TYPE, and PRINT statements, other I/O record-related statements are limited to a specific file organization. For instance:

- The DELETE statement only applies to relative files. (Detecting deleted records is only available if the `-vms` option was specified when the program was compiled.)
- The BACKSPACE statement only applies to sequential files open for sequential access.
- The REWIND statement only applies to sequential files open for sequential access and to direct access files.
- The ENDFILE statement only applies to certain types of sequential files open for sequential access and to direct access files.
- The UNLOCK statement only applies to records in relative or sequential files.

The file-related statements (OPEN, INQUIRE, and CLOSE) apply to any relative or sequential file.

# Forms of I/O Statements

Each type of record I/O statement can be coded in a variety of forms. The form you select depends on the nature of your data and how you want it treated. When opening a file, specify the form using the FORM specifier.

The following are the forms of I/O statements:

- *Formatted I/O statements* contain explicit format specifiers that are used to control the translation of data from internal (binary) form within a program to external (readable character) form in the records, or vice versa.
- *List-directed* and *namelist I/O statements* are similar to formatted statements in function. However, they use different mechanisms to control the translation of data: formatted I/O statements use explicit format specifiers, and list-directed and namelist I/O statements use data types.

- *Unformatted I/O statements* do not contain format specifiers and therefore do not translate the data being transferred (important when writing data that will be read later).

Formatted, list-directed, and namelist I/O forms require translation of data from internal (binary) form within a program to external (readable character) form in the records. Consider using unformatted I/O for the following reasons:

- Unformatted data avoids the translation process, so I/O tends to be faster.
- Unformatted data avoids the loss of precision in floating-point numbers when the output data will subsequently be used as input data.
- Unformatted data conserves file storage space (stored in binary form).

To write data to a file using formatted, list-directed, or namelist I/O statements, specify FORM= 'FORMATTED' when opening the file. To write data to a file using unformatted I/O statements, specify FORM= 'UNFORMATTED' when opening the file.

Data written using formatted, list-directed, or namelist I/O statements is referred to as formatted data; data written using unformatted I/O statements is referred to as unformatted data.

When reading data from a file, you should use the same I/O statement form that was used to write the data to the file. For instance, if data was written to a file with a formatted I/O statement, you should read data from that file with a formatted I/O statement.

Although I/O statement form is usually the same for reading and writing data in a file, a program can read a file containing unformatted data (using unformatted input) and write it to a separate file containing formatted data (using formatted output). Similarly, a program can read a file containing formatted data and write it to a different file containing unformatted data.

You can access records in any sequential or relative file using sequential access. For relative files and certain (fixed-length) sequential files, you can also access records using direct access.

The table below shows the main record I/O statements, by category, that can be used in Intel Fortran programs.

| File Type, Access, and I/O Form | Available Statements |
|---|---|
| **External file, sequential access** | |
| Formatted | READ, WRITE, PRINT, ACCEPT, TYPE, REWRITE |

| List-directed | READ, WRITE, PRINT, ACCEPT, TYPE |
|---|---|
| Namelist | READ, WRITE, PRINT, ACCEPT, TYPE |
| Unformatted | READ, WRITE, REWRITE |
| **External file, direct access** | |
| Formatted | READ, WRITE, REWRITE |
| Unformatted | READ, WRITE, REWRITE |
| **Internal file** | |
| Formatted | READ, WRITE |
| List-directed | READ, WRITE |
| Unformatted | None |

**Note**

You can use the REWRITE statement only for relative files, using direct access.

**Files and File Characteristics**

# Files and File Characteristics Overview

See these topics:

File Organization

Internal Files and Scratch Files

Record Types

Record Overhead

Record Length

# File Organization

*File organization* refers to the way records are physically arranged on a storage device.

Intel Fortran supports two kinds of file organization:

- Sequential

- Relative

The default file organization is always ORGANIZATION= 'SEQUENTIAL' for an OPEN statement. The organization of a file is specified by means of the ORGANIZATION specifier in the OPEN statement.

You can store sequential files on magnetic tape or disk devices, and can use other peripheral devices, such as terminals, pipes, and line printers as sequential files.

You must store relative files on a disk device.

## Sequential Organization

A sequentially organized file consists of records arranged in the sequence in which they are written to the file (the first record written is the first record in the file, the second record written is the second record in the file, and so on). As a result, records can be added only at the end of the file.

Sequential files are usually read sequentially, starting with the first record in the file. Sequential files with a fixed-length record type that are stored on disk can also be accessed by relative record number (direct access).

## Relative Organization

Within a relative file are numbered positions, called cells. These cells are of fixed equal length and are consecutively numbered from 1 to *n*, where 1 is the first cell, and *n* is the last available cell in the file. Each cell either contains a single record or is empty.

Records in a relative file are accessed according to *cell number*. A cell number is a record's relative record number (its location relative to the beginning of the file). By specifying relative record numbers, you can directly retrieve, add, or delete records regardless of their locations (direct access). (Detecting deleted records is only available if you specified the -vms option when the program was compiled.)

When creating a relative file, use the RECL value to determine the size of the fixed-length cells. Within the cells, you can store records of varying length, as long as their size does not exceed the cell size.

# Internal Files and Scratch Files

Intel Fortran also supports two other types of files that are not file organizations:

- Internal files
- Scratch files

# Internal Files

When you use sequential access, you can use an internal file to reference character data in a buffer. The transfer occurs between internal storage and internal storage (unlike external files), such as between character variables and a character array.

An internal file consists of any of the following:

- Character variable
- Character-array element
- Character array
- Character substring
- Character array section without a vector subscript

Instead of specifying a unit number for the READ or WRITE statement, use an internal file specifier in the form of a character scalar memory reference or a character-array name reference.

An internal file is a designated internal storage space (variable buffer) of characters that is treated as a sequential file of fixed-length records. To perform internal I/O, use formatted and list-directed sequential READ and WRITE statements. You cannot use file-related statements such as OPEN and INQUIRE on an internal file (no unit number is used).

If an internal file is made up of a single character variable, array element, or substring, that file comprises a single record whose length is the same as the length of the character variable, array element, or substring it contains. If an internal file is made up of a character array, that file comprises a sequence of records, with each record consisting of a single array element. The sequence of records in an internal file is determined by the order of subscript progression.

A record in an internal file can be read only if the character variable, array element, or substring comprising the record has been defined (a value has been assigned to the record).

Prior to each READ and WRITE statement, an internal file is always positioned at the beginning of the first record.

# Scratch Files

Scratch files are created by specifying STATUS= ' SCRATCH ' in an OPEN
statement. By default, these temporary files are created in (and later deleted
from) the directory specified in the OPEN statement DEFAULTFILE (if specified).

# Record Types

*Record type* refers to whether records stored in a file are all the same length, are
of varying length, or use other conventions to define where one record ends and
another begins.

You can use fixed-length and variable-length record types with sequential or
relative files. You can use any of the record types with sequential files. Relative
files require the fixed-length record type.

When creating a new file or opening an existing file, specify one of the record
types described below.

See also Format of Record Types,

## Fixed-Length Record Type

Records in a file must be the same length.

You must specify the record length (RECL) when the file is opened.

See also Fixed-Length Records.

## Variable-Length Record Type

Records in a file can vary in length.

Record length information is stored as control bytes at the beginning and end of
each record.

See also Variable-Length Records.

## Segmented Record Type

This pertains to a single logical record containing one or more unformatted
records of varying length, which can only be used for unformatted sequential
access.

Avoid the segmented record type when the application requires that the same file be used for programs written in languages other than Fortran and for non-Intel platforms.

See also Sesgmented Records.

# Stream Record Type

A stream file is not grouped into records and uses no record delimiters.

Stream files contain character or binary data that is read or written to the extent of the variables specified. Specify CARRIAGECONTROL= ' NONE ' for stream files.

See also Stream Files.

# Stream_LF and Stream_CR Record Type

Records are of varying length where the line feed (LF) or the carriage return (CR) character serve as record delimiters (LF for Stream_LF files and CR for Stream_CR files).

Stream_LF files must not contain embedded LF characters or use CARRIAGECONTROL= ' LIST ' . Instead, specify CARRIAGECONTROL= ' NONE ' . Stream_CR files must not contain embedded CR characters. The Stream_LF record type is the usual record type for text files.

See also Stream_LF and Stream_CR Records.

# Choosing a Record Type

Before you choose a record type, consider whether your application will use formatted or unformatted data. If you are using formatted data, you can choose any record type except segmented. If you are using unformatted data, avoid the Stream, Stream_CR, and Stream_LF record types.

The segmented record type can only be used for unformatted sequential access with sequential files. You should not use segmented records for files that are read by programs written in languages other than Intel Fortran.

The Stream, Stream_CR, Stream_LF, and segmented record types can be used only with sequential files.

The default record type (RECORDTYPE) depends on the values for the ACCESS and FORM specifiers for the OPEN statement.

The record type of the file is not maintained as an attribute of the file. The results of using a record type other than the one used to create the file are indeterminate.

An I/O record is a collection of fields (data items) that are logically related and are usually processed as a unit.

Unless you specify nonadvancing I/O (ADVANCE specifier), each Intel Fortran I/O statement transfers at least one record.

# Record Overhead

*Record overhead* refers to bytes associated with each record that are used internally by the file system and are not available when a record is read or written. Knowing the record overhead helps when estimating the storage requirements for an application. Although the overhead bytes exist on the storage media, do not include them when specifying the record length with the RECL specifier in an OPEN statement.

The various record types each require a different number of bytes for record overhead, as described in the table below:

| Record Type | File Organization | Record Overhead |
| --- | --- | --- |
| Fixed-length | Sequential | None. |
| Fixed-length | Relative | None if the `-vms` option was omitted (the default). One byte if the `-vms` option was specified. |
| Variable-length | Sequential | Eight bytes per record. |
| Segmented | Sequential | Four bytes per record. One additional padding byte (space) is added if the specified record size is an odd number. |
| Stream | Sequential | None required. |
| Stream_CR | Sequential | One byte per record. |
| Stream_LF | Sequential | One byte per record. |

# Record Length

Use the RECL specifier to specify the record length.

The units used for specifying record length depend on the form of the data:

- Formatted files (FORM= ' FORMATTED ' ): Specify the record length in bytes.
- Unformatted files (FORM= ' UNFORMATTED ' ): Specify the record length in 4-byte units, unless you specify the `-assume byterecl` option to request 1-byte units.

For all but variable-length sequential records on 64-bit addressable systems, the maximum record length is 2.147 billion bytes (2,147,483,647 minus the bytes for record overhead). For variable-length sequential records on 64-bit addressable systems, the theoretical maximum record length is about 17,000 gigabytes. When considering very large record sizes, also consider limiting factors such as system virtual memory.

# Accessing and Assigning Files

Most I/O operations involve a disk file, keyboard, or screen display. Other devices can also be used:

- Sockets can be read from or written to if a USEROPEN routine (usually written in C) is used to open the socket.
- Pipes opened for read and write access block (wait until data is available) if you issue a READ to an empty pipe.
- Pipes opened for read-only access return EOF if you issue a READ to an empty pipe.

You can access the terminal screen or keyboard by using preconnected files.

## Assigning Files to Logical Units

You can choose to assign files to logical units by using one of the following methods:

- Using default values, such as a preconnected unit
- Supplying a file name (and possibly a directory) in an OPEN statement
- Using environment variables

### Using Default Values

In the following example, the PRINT statement is associated with a preconnected unit (`stdout`) by default.

```
PRINT *,100
```

The READ statement associates the logical unit 7 with the file `fort.7` (because the FILE specifier was omitted) by default:

```
OPEN (UNIT=7,STATUS='NEW')
READ (7,100)
```

### Supplying a File Name in an OPEN Statement

For example:

```
OPEN (UNIT=7, FILE='FILNAM.DAT', STATUS='OLD')
```

The FILE specifier in an OPEN statement typically specifies only a file name (such as `testdata`) or contains both a directory and file name (such as `/usr/proj/testdata`).

The DEFAULTFILE specifier in an OPEN statement typically specifies a pathname that contains only a directory (such as `/usr/proj/`) or both a directory and file name (such as `/usr/proj/testdata`).

### Implied OPEN

Performing an implied OPEN means that the FILE and DEFAULTFILE specifier values are not specified and an environment variable is used, if present. Thus, if you used an implied OPEN, or if the FILE specifier in an OPEN statement did not specify a file name, you can use an environment variable to specify a file name or a pathname that contains both a directory and file name.

### Using Environment Variables

You can use shell commands to set the appropriate environment variable to a value that indicates a directory (if needed) and a file name to associate a unit with an external file.

Intel Fortran recognizes environment variables for each logical I/O unit number in the form of `FORT`$n$, where $n$ is the logical I/O unit number. If a file name is not specified in the OPEN statement and the corresponding `FORT`$n$ environment variable is not set for that unit number, Intel Fortran generates a file name in the form `fort.`$n$ , where $n$ is the logical unit number.

# Implied Intel Fortran Logical Unit Numbers

The ACCEPT, PRINT, and TYPE statements, and the use of an asterisk (*) in place of a unit number in READ and WRITE statements, do not include an explicit logical unit number.

Each of these Fortran statements uses an implicit internal logical unit number and environment variable. Each environment variable is in turn associated by default with one of the Fortran file names that are associated with standard I/O files. The table below shows these relationships:

| Intel Fortran statement | Environment variable when –`vms` specified | Environment variable when –`vms` omitted | Standard I/O file name |
|---|---|---|---|
| READ (*,f) iolist | `FOR_READ` | `FORT5` | `stdin` |
| READ f,iolist | `FOR_READ` | `FORT5` | `stdin` |
| ACCEPT f,iolist | `FOR_ACCEPT` | `FORT5` | `stdin` |
| WRITE (*,f) iolist | `FOR_PRINT` | `FORT6` | `stdout` |
| PRINT f,iolist | `FOR_PRINT` | `FORT6` | `stdout` |
| TYPE f,iolist | `FOR_TYPE` | `FORT6` | `stdout` |

You can change the file associated with these Intel Fortran environment variables, as you would any other environment variable, by means of the environment variable assignment command. For example:

```
setenv FOR_READ /usr/users/smith/test.dat
```

After executing the preceding command, the environment variable for the READ statement using an asterisk refers to file `test.dat` in directory `/usr/users/smith`.

# Default Pathnames and File Names

Intel Fortran provides the following possible ways of specifying all or part of a file specification (directory and file name), such as `/usr/proj/testdata`:

- The FILE specifier in an OPEN statement typically specifies only a file name (such as `testdata`) or contains both a directory and file name (such as `/usr/proj/testdata`).
- The DEFAULTFILE specifier in an OPEN statement typically specifies a pathname that contains only a directory (such as `/usr/proj/`) or both a directory and file name (such as `/usr/proj/testdata`).
- If you used an implied OPEN or if the FILE specifier in an OPEN statement did not specify a file name, you can use an environment variable to specify a file name or a pathname that contains both a directory and file name.

# Examples of Applying Default Pathnames and File Names

For example, for an implied OPEN of unit number 3, Intel Fortran would check the environment variable FORT3. If the environment variable FORT3 was set, its value is used. If it is not set, the system supplies the file name `fort.3`.

In the following table, assume the current directory is `/usr/smith` and the I/O uses unit 1, as in the statement READ (1,100).

| OPEN FILE value | OPEN DEFAULTFILE value | FORT1 environment variable value | Resulting pathname |
|---|---|---|---|
| not specified | not specified | not specified | `/usr/smith/fort.1` |
| not specified | not specified | `test.dat` | `/usr/smith/test.dat` |
| not specified | not checked | `/usr/tmp/t.dat` | `/usr/tmp/t.dat` |
| not specified | `/tmp` | not specified | `/tmp/fort.1` |
| not specified | `/tmp` | `testdata` | `/tmp/testdata` |
| not specified | `/usr` | `lib/testdata` | `/usr/lib/testdata` |
| `file.dat` | `/usr/group` | not checked | `/usr/group/file.dat` |
| `/tmp/file.dat` | not checked | not checked | `/tmp/file.dat` |
| `file.dat` | not specified | not specified | `/usr/smith/file.dat` |

When the resulting file pathname begins with a tilde character (~), C-shell-style pathname substitution is used (regardless of what shell is being used), such as a top-level directory (below the root). For additional information on tilde pathname substitution, see `csh(1)`.

# Rules for Applying Default Pathnames and File Names

Intel Fortran determines file name and the directory path based on certain rules. It determines a file name string as follows:

- If the FILE specifier is present, its value is used.
- If the FILE specifier is not present, Intel Fortran examines the corresponding environment variable. If the corresponding environment variable is set, that value is used. If the corresponding environment variable is not set, a file name in the form `fort.`*n* is used.

Once Intel Fortran determines the resulting file name string, it determines the directory (which optionally precedes the file name) as follows:

- If the resulting file name string contains an absolute pathname, it is used and the DEFAULTFILE specifier, environment variable, and current directory values are ignored.
- If the resulting file name string does not contain an absolute pathname, Intel Fortran examines the DEFAULTFILE specifier and current directory value: If the corresponding environment variable is set and specifies an absolute pathname, that value is used. Otherwise, the DEFAULTFILE specifier value, if present, is used. If the DEFAULTFILE specifier is not present, Intel Fortran uses the current directory as an absolute pathname.

# Using Preconnected Standard I/O Files

If you do not use an OPEN statement to open logical unit 5, 6, or 0 and do not set the appropriate environment variable (FORT*n*), Intel Fortran at run time implicitly opens (preconnected) units 5, 6, and 0 and associates them with the following operating system standard I/O files:

| Unit | Environment Variable | Equivalent Linux* Standard I/O File |
|------|---------------------|-------------------------------------|
| 5 | FORT5 | Standard input, stdin |
| 6 | FORT6 | Standard output, stdout |
| 0 | FORT0 | Standard error, stderr |

You can change these preconnected files by doing one of the following:

- Using an OPEN statement to open unit 5, 6, or 0. When you explicitly OPEN a file for unit 5, 6, or 0, the OPEN statement keywords specify the file-related information to be used instead of the preconnected standard I/O file.
- Setting the appropriate environment variable (FORT*n*) to redirect I/O to an external file.

To redirect input or output from the standard preconnected files at run time, you can set the appropriate environment variable or use the appropriate shell redirection character in a pipe (such as > or <).

# Opening Files: OPEN Statement

To open a file, you should use a preconnected file (such as for terminal output) or explicitly open files with an OPEN statement. Although you can also implicitly open a file, this prevents you from using the OPEN statement to specify the file connection characteristics and other information.

# OPEN Statement Specifiers

The OPEN statement connects a unit number with an external file and allows you to explicitly specify file attributes and run-time options using OPEN statement specifiers. Once you open a file, you should close it before opening it again unless it is a preconnected file.

If you open a unit number that was opened previously (without being closed), one of the following occurs:

- If you specify a file specification that does not match the one specified for the original open, the Intel Fortran run-time system closes the file and then reopens it. This resets the current record position for the second file.
- If you specify a file specification that matches the one specified for the original open, the file is reconnected without the internal equivalent of the CLOSE and OPEN. This lets you change one or more OPEN statement run-time specifiers while maintaining the record position context.

You can use the INQUIRE statement to obtain information about whether or not a file is opened by your program.

Especially when creating a new file using the OPEN statement, examine the defaults (see the description of the OPEN statement in the *Intel Fortran Language Reference Manual*) or explicitly specify file attributes with the appropriate OPEN statement specifiers.

## Specifiers for File and Unit Information

These specifiers identify file and unit information:

- UNIT specifies the logical unit number.
- FILE (or NAME) and DEFAULTFILE specify the directory and/or file name of an external file.
- STATUS or TYPE indicates whether to create a new file, overwrite an existing file, open an existing file, or use a scratch file.
- STATUS or DISPOSE specifies the file existence status after CLOSE.

## Specifiers for File and Record Characteristics

These specifiers identify file and record characteristics:

- ORGANIZATION indicates the file organization (sequential or relative).
- RECORDTYPE indicates which record type to use.
- FORM indicates whether records are formatted or unformatted.
- CARRIAGECONTROL indicates the terminal control type.

- RECL or RECORDSIZE specifies the record size.

## Specifier for Special File Open Routine

USEROPEN names the routine that will open the file to establish special context that changes the effect of subsequent Intel Fortran I/O statements.

## Specifiers for File Access, Processing, and Position

These specifiers identify file access, processing, and position:

- ACCESS indicates the access mode (direct or sequential).
- SHARED indicates that other users can access the same file and activates record locking. Ignored in the current version of Intel Fortran.
- POSITION indicates whether to position the file at the beginning of file, before the end-of-file record, or leave it as is (unchanged).
- ACTION or READONLY indicates whether statements will be used to only read records, only write records, or both read and write records.
- MAXREC specifies the maximum record number for direct access.
- ASSOCIATEVARIABLE specifies the variable containing the next record number for direct access.

## Specifiers for Record Transfer Characteristics

These specifiers identify record transfer characteristics:

- BLANK indicates whether to ignore blanks in numeric fields.
- DELIM specifies the delimiter character for character constants in list-directed or namelist output.
- PAD, when reading formatted records, indicates whether padding characters should be added if the item list and format specification require more data than the record contains.
- BLOCKSIZE specifies the block physical I/O buffer size.
- BUFFERCOUNT specifies the number of physical I/O buffers.
- CONVERT specifies the format of unformatted numeric data.

## Specifiers for Error-Handling Capabilities

These specifiers are used for error handling:

- ERR specifies a label to branch to if an error occurs.
- IOSTAT specifies the integer variable to receive the error (IOSTAT) number if an error occurs.

## Specifier for File Close Action

DISPOSE identifies the action to take when the file is closed.

## Coding File Locations in an OPEN Statement

You can use the FILE and DEFAULTFILE specifiers of the OPEN statement to specify the complete definition of a particular file to be opened on a logical unit. (The *Language Reference Manual* describes the OPEN statement in greater detail.)

For example:

```
OPEN (UNIT=4, FILE='/usr/users/smith/test.dat',
STATUS='OLD')
```

The file `test.dat` in directory `/usr/users/smith` is opened on logical unit 4. No defaults are applied, because both the directory and file name were specified. The value of the FILE specifier can be a character constant, variable, or expression.

In the following interactive example, the user supplies the file name and the DEFAULTFILE specifier supplies the default values for the full pathname string. The file to be opened is in `/usr/users/smith` and is concatenated with the file name typed by the user into the variable DOC:

```
CHARACTER(LEN=9) DOC
WRITE (6,*)  'Type file name '
READ (5,*) DOC
OPEN (UNIT=2, FILE=DOC,
DEFAULTFILE='/usr/users/smith',STATUS='OLD')
```

A slash is appended to the end of the default file string if it does not have one.

# Obtaining File Information: INQUIRE Statement

The INQUIRE statement returns information about a file and has three forms:

- Inquiry by unit
- Inquiry by file name
- Inquiry by output item list

## Inquiry by Unit

An inquiry by unit is usually done for an opened (connected) file. An inquiry by unit causes the Intel Fortran RTL to check whether the specified unit is connected or not. One of the following occurs, depending on whether the unit is connected or not:

If the unit is connected:

- The EXIST and OPENED specifier variables indicate a true value.
- The pathname and file name are returned in the NAME specifier variable (if the file is named).
- Other information requested on the previously connected file is returned.
- Default values are usually returned for the INQUIRE specifiers also associated with the OPEN statement.
- The RECL value unit for connected formatted files is always 1-byte units. For unformatted files, the RECL unit is 4-byte units, unless you specify the `-assume byterecl` option to request 1-byte units.

If the unit is not connected:

- The OPENED specifier indicates a false value.
- The unit NUMBER specifier variable is returned as a value of -1.
- Any other information returned will be undefined or default values for the various specifiers.0

For example, the following INQUIRE statement shows whether unit 3 has a file connected (OPENED specifier) in logical variable I_OPENED, the name (case-sensitive) in character variable I_NAME, and whether the file is opened for READ, WRITE, or READWRITE access in character variable I_ACTION:

```
INQUIRE (3, OPENED=I_OPENED, NAME=I_NAME, ACTION=I_ACTION)
```

# Inquiry by File Name

An inquiry by name causes the Intel Fortran RTL to scan its list of open files for a matching file name. One of the following occurs, depending on whether a match occurs or not:

If a match occurs:

- The EXIST and OPENED specifier variables indicate a true value.
- The pathname and file name are returned in the NAME specifier variable.
- The UNIT number is returned in the NUMBER specifier variable.
- Other information requested on the previously connected file is returned.
- Default values are usually returned for the INQUIRE specifiers also associated with the OPEN statement.

- The RECL value unit for connected formatted files is always 1-byte units. For unformatted files, the RECL unit is 4-byte units, unless you specify the `-assume byterecl` option to request 1-byte units.

If no match occurs:

- The OPENED specifier variable indicates a false value.
- The unit NUMBER specifier variable is returned as a value of -1.
- The EXIST specifier variable indicates (true or false) whether the named file exists on the device or not.
- If the file does exist, the NAME specifier variable contains the pathname and file name.
- Any other information returned will be default values for the various specifiers, based on any information specified when calling INQUIRE.

The following INQUIRE statement returns whether the file named `log_file` is connected in logical variable I_OPEN, whether the file exists in logical variable I_EXIST, and the unit number in integer variable I_NUMBER:

```
INQUIRE (FILE='log_file', OPENED=I_OPEN, EXIST=I_EXIST,
NUMBER=I_NUMBER)
```

## Inquiry by Output Item List

Unlike inquiry by unit or inquiry by name, inquiry by output item list does not attempt to access any external file. It returns the length of a record for a list of variables that would be used for unformatted WRITE, READ, and REWRITE statements. The following INQUIRE statement returns the maximum record length of the variable list in variable I_RECLENGTH. This variable is then used to specify the RECL value in the OPEN statement:

```
INQUIRE (IOLENGTH=I_RECLENGTH) A, B, H
OPEN (FILE='test.dat', FORM='UNFORMATTED', RECL=I_RECLENGTH,
UNIT=9)
```

For an unformatted file, the RECL value is returned using 4-byte units, unless you specify the `-assume byterecl` option to request 1-byte units.

# Closing a File: CLOSE Statement

Usually, any external file opened should be closed by the same program before it completes. The CLOSE statement disconnects the unit and its external file. You must specify the unit number (UNIT specifier) to be closed.

You can also specify:

- Whether the file should be deleted or kept (STATUS specifier)
- Error handling information (ERR and IOSTAT specifiers)

To delete a file when closing it:

- In the OPEN statement, specify the ACTION keyword (such as ACTION='READ'). Avoid using the READONLY keyword, because a file opened using the READONLY keyword cannot be deleted when it is closed.
- In the CLOSE statement, specify the keyword STATUS='DELETE'.

If you opened an external file and did an inquire by unit, but do not like the default value for the ACCESS specifier, you can close the file and then reopen it, explicitly specifying the ACCESS desired.

There usually is no need to close preconnected units. Internal files are neither opened nor closed.

**Record Operations**

# Record Operations Overview

See these topics:

Record I/O Statement Specifiers

Record Access

File Sharing

Specifying the Initial Record Position

Advancing and Nonadvancing Record I/O

Record Transfer

# Record I/O Statement Specifiers

After you open a file or use a preconnected file, you can use the following statements:

- READ, WRITE, ACCEPT, and PRINT to perform record I/O.
- BACKSPACE, ENDFILE, and REWIND to set record position within the file.
- DELETE, REWRITE, TYPE, and FIND to perform various operations.

The record I/O statement must use the appropriate record I/O form (formatted, list-directed, namelist, or unformatted).

You can use the following specifiers with the READ and WRITE record I/O statements:

- UNIT specifies the unit number to or from which input or output will occur.
- END specifies a label to branch to if end-of-file occurs; only applies to input statements on sequential files.
- ERR specifies a label to branch to if an error occurs.
- IOSTAT specifies an integer variable to contain the error number if an error occurs.
- FMT specifies a label of a FORMAT statement or character data specifying a FORMAT.
- NML specifies the name of a NAMELIST.
- REC specifies a record number for direct access.

When using nonadvancing I/O, use the ADVANCE, EOR, and SIZE specifiers.

When using the REWRITE statement, you can use the UNIT, FMT, ERR, and IOSTAT specifiers.

# Record Access

*Record access* refers to how records will be read from or written to a file, regardless of the file's organization. Record access is specified each time you open a file; it can be different each time. The type of record access permitted is determined by the combination of file organization and record type.

For instance, you can:

- Add records to a sequential file with ORGANIZATION= ' SEQUENTIAL ' and POSITION= ' APPEND ' (or use ACCESS= ' APPEND ' ).
- Add records sequentially by using multiple WRITE statements, close the file, and then open it again with ORGANIZATION= ' SEQUENTIAL ' and ACCESS= ' SEQUENTIAL ' (or ACCESS= ' DIRECT ' if the sequential file has fixed-length records).

## Sequential Access

Sequential access transfers records sequentially to or from files or I/O devices such as terminals. You can use sequential I/O with any type of supported file organization and record type.

If you select sequential access mode for files with sequential or relative organization, records are written to or read from the file starting at the beginning of the file and continuing through it, one record after another. A particular record can be retrieved only after all of the records preceding it have been read; new records can be written only at the end of the file.

# Direct Access

Direct access transfers records selected by record number to and from either sequential files stored on disk with a fixed-length record type or relative organization files.

If you select direct access mode, you can determine the order in which records are read or written. Each READ or WRITE statement must include the relative record number, indicating the record to be read or written.

You can directly access a sequential disk file only if it contains fixed-length records. Because direct access uses cell numbers to find records, you can enter successive READ or WRITE statements requesting records that either precede or follow previously requested records. For example, the first of the following statements reads record 24; the second reads record 10:

```
READ (12,REC=24) I
READ (12,REC=10) J
```

# Limitations of Record Access by File Organization and Record Type

You can use both access modes on sequential and relative files. However, direct access to a sequential organization file can only be done if the file resides on disk and contains fixed-length records.

The table below summarizes the types of access permitted for the various combinations of file organizations and record types.

| Record Type | Sequential Access? | Direct Access? |
|---|---|---|
| **Sequential file organization** | | |
| Fixed | Yes | Yes |
| Variable | Yes | No |
| Segmented | Yes | No |

| Stream | Yes | No |
|---|---|---|
| Stream_CR | Yes | No |
| Stream_LF | Yes | No |
| **Relative file organization** | | |
| Fixed | Yes | Yes |

📓**Note**

Direct access and relative files require that the file resides on a disk device.

# File Sharing

Depending on the value specified by the ACTION (or READONLY) specifier in the OPEN statement, the file will be opened by your program for reading, writing, or both reading and writing records. This simply checks that the program itself executes the type of statements intended.

For performance reasons, record locking and shared-file checking are not supported by Intel Fortran. When you open the file, access is always granted, regardless of whether:

- The OPEN statement SHARED specifier was specified
- Other processes have already opened the file

You might open a file for writing records (or reading and writing records) and know another process might simultaneously have the file open and be writing records. In this case, you need to coordinate access times among those processes to handle the possibility of simultaneous WRITE and REWRITE statements on the same record positions.

# Specifying the Initial Record Position

When you open a disk file, you can use the OPEN statement POSITION specifier to request one of the following initial record positions within the file:

- The initial position before the first record (POSITION='REWIND'). A sequential access READ or WRITE statement will read or write the first record in the file.
- A point beyond the last record in the file (POSITION='APPEND'), just before the end-of-file record, if one exists. For a new file, this is the initial position before the first record (same as 'REWIND'). You might specify 'APPEND' before you write records to an existing sequential file using sequential access.

- The current position (POSITION='ASIS'). This is usually used only to maintain the current record position when reconnecting a file. The second OPEN specifies the same unit number and specifies the same file name (or omits it), which leaves the file open, retaining the current record position.

However, if the second OPEN specifies a different file name for the same unit number, the current file will be closed and the different file will be opened.

The following I/O statements allow you to change the current record position:

- REWIND sets the record position to the initial position before the first record. A sequential access READ or WRITE statement would read or write the first record in the file.
- BACKSPACE sets the record position to the previous record in a file. Using sequential access, if you wrote record 5, issued a BACKSPACE to that unit, and then read from that unit, you would read record 5.
- ENDFILE writes an end-of-file marker. This is typically done after writing records using sequential access just before you close the file.

Unless you use nonadvancing I/O, reading and writing records usually advances the current record position by one record. More than one record might be transferred using a single record I/O statement.

# Advancing and Nonadvancing Record I/O

After you open a file, if you omit the ADVANCE specifier (or specify ADVANCE= 'YES' ) in READ and WRITE statements, advancing I/O (normal FORTRAN-77 I/O) will be used for record access. When using advancing I/O:

- Record I/O statements transfer one entire record (or multiple records).
- Record I/O statements advance the current record position to a position before the next record.

You can request nonadvancing I/O for the file by specifying the ADVANCE= ' NO ' specifier in a READ and WRITE statement. You can use nonadvancing I/O only for sequential access to external files using formatted I/O (not list-directed or namelist).

When you use nonadvancing I/O, the current record position does not change, and part of the record might be transferred, unlike advancing I/O where one entire record or records are always transferred.

You can alternate between advancing and nonadvancing I/O by specifying different values for the ADVANCE specifier ( ' YES ' and ' NO ' ) in the READ and WRITE record I/O statements.

When reading records with either advancing or nonadvancing I/O, you can use the END specifier to branch to a specified label when the end of the file is read.

Because nonadvancing I/O might not read an entire record, it also supports an EOR specifier to branch to a specified label when the end of the record is read. If you omit the EOR and the IOSTAT specifiers when using nonadvancing I/O, an error results when the end-of-record is read.

When using nonadvancing input, you can use the SIZE specifier to return the number of characters read. For example, in the following READ statement, SIZE=X (where variable X is an integer) returns the number of characters read in X and an end-of-record condition causes a branch to label 700:

```
150 FORMAT (F10.2, F10.2, I6)
    READ (UNIT=20, FMT=150, SIZE=X, ADVANCE='NO', EOR=700)
A, F, I
```

# Record Transfer

I/O statements transfer all data as records. The amount of data that a record can contain depends on the following circumstances:

- With formatted I/O (except for fixed-length records), the number of items in the I/O statement and its associated format specifier jointly determine the amount of data to be transferred.
- With namelist and list-directed output, the items listed in the NAMELIST statement or I/O statement list (in conjunction with the NAMELIST or list-directed formatting rules) determine the amount of data to be transferred.
- With unformatted I/O (except for fixed-length records), the I/O statement alone specifies the amount of data to be transferred.
- When you specify fixed-length records (RECORDTYPE= 'FIXED'), all records are the same size. If the size of an I/O record being written is less than the record length (RECL), extra bytes are added (padding).

Typically, the data transferred by an I/O statement is read from or written to a single record. It is possible, however, for a single I/O statement to transfer data from or to more than one record, depending on the form of I/O used.

## Input Record Transfer

When using advancing I/O, if an input statement specifies fewer data fields (less data) than the record contains, the remaining fields are ignored.

If an input statement specifies more data fields than the record contains, one of the following occurs:

- For formatted input using advancing I/O, if the file was opened with PAD='YES', additional fields are read as spaces. If the file is opened with PAD='NO', an error occurs (the input statement should not specify more data fields than the record contains).
- For formatted input using nonadvancing I/O (ADVANCE='NO'), an end-of-record (EOR) condition is returned. If the file was opened with PAD='YES', additional fields are read as spaces.
- For list-directed input, another record is read.
- For NAMELIST input, another record is read.
- For unformatted input, an error occurs.

# Output Record Transfer

If an output statement specifies fewer data fields than the record contains (less data than required to fill a record), the following occurs:

- With fixed-length records (RECORDTYPE= ' FIXED ' ), all records are the same size. If the size of an I/O record being written is less than the record length (RECL), extra bytes are added (padding) in the form of spaces (for a formatted record) or zeros (for an unformatted record).
- With other record types, the fields present are written and those omitted are not written (might result in a short record).

If the output statement specifies more data than the record can contain, an error occurs, as follows:

- With formatted or unformatted output using fixed-length records, if the items in the output statement and its associated format specifier result in a number of bytes that exceeds the maximum record length (RECL), an error occurs.
- With formatted or unformatted output not using fixed-length records, if the items in the output statement and its associated format specifier result in a number of bytes that exceeds the maximum record length (RECL), the Intel Fortran RTL attempts to increase the RECL value and write the longer record. To obtain the RECL value, use an INQUIRE statement.
- For list-directed output and namelist output, if the data specified exceeds the maximum record length (RECL), another record is written.

# User-Supplied OPEN Procedures: USEROPEN Specifier

You can use the USEROPEN specifier in an Intel Fortran OPEN statement to pass control to a routine that directly opens a file. The called routine can use system calls or library routines to open the file and establish special context that changes the effect of subsequent Intel Fortran I/O statements.

The Intel Fortran RTL I/O support routines call the USEROPEN function in place of the system calls usually used when the file is first opened for I/O. The USEROPEN specifier in an OPEN statement specifies the name of a function to receive control. The called function must open the file (or pipe) and return the file descriptor of the file when it returns control to the RTL.

When opening the file, the called function usually specifies options different from those provided by a normal OPEN statement.

You can obtain the file descriptor from the Intel Fortran RTL for a specific unit number with the `getfd` routine.

Although the called function can be written in other languages (such as Fortran), C is usually the best choice for making system calls, such as open or create.

## Syntax and Behavior of the USEROPEN Specifier

The USEROPEN specifier for the OPEN statement has the form:

```
USEROPEN = function-name
```

*function-name* represents the name of an external function. In the calling program, the function must be declared in an EXTERNAL statement. For example, the following Intel Fortran code might be used to call the USEROPEN procedure UOPEN (known to the linker as `uopen_`):

```
EXTERNAL   UOPEN
INTEGER    UOPEN
.
.
.
OPEN (UNIT=10, FILE='/usr/test/data', STATUS='NEW',
USEROPEN=UOPEN)
```

During the OPEN statement, the `uopen_` function receives control. The function opens the file, may perform other operations, and subsequently returns control (with the file descriptor) to the RTL.

If the USEROPEN function is written in C, declare it as a C function that returns a 4-byte integer (`int`) result to contain the file descriptor. For example:

```
int    uopen_ (                     (1)
       char  *file_name,            (2)
       int   *open_flags,           (3)
       int   *create_mode,          (4)
       int   *lun,                  (5)
       int   file_length);          (6)
```

The function definition and the arguments passed from the Intel Fortran RTL are as follows:

1. The function must be declared as a 4-byte integer (`int`).
2. The first argument is the pathname (includes the file name) to be opened.
3. The open flags are described in the header file `/usr/include/sys/file.h` or `open(2)`.
4. The create mode (protection needed when creating a file) is described in `open(2)`.
5. The fourth argument is the logical unit number.
6. The fifth (last) argument is the pathname length (hidden length argument of the pathname).

Of the arguments, the `open` system call (see `open(2)`) requires the passed pathname, the open flags (that define the type access needed, whether the file exists, and so on), and the create mode. The logical unit number specified in the OPEN statement is passed in case the USEROPEN function needs it. The hidden length of the pathname is also passed.

When creating a new file, the create system call might be used in place of open (see `create(2)`). You can usually use other appropriate system calls or library routines within the USEROPEN function.

In most cases, the USEROPEN function modifies the open flags argument passed by the Intel Fortran RTL or uses a new value before the open (or create) system call. After the function opens the file, it must return control to the RTL.

If the USEROPEN function is written in Fortran, declare it as a FUNCTION with an INTEGER (KIND=4) result, perhaps with an interface block. In any case, the called function must return the file descriptor as a 4-byte integer to the RTL.

If your application requires that you use C to perform the file open and close, as well as all record operations, call the appropriate C procedure from the Intel Fortran program without using the Fortran OPEN statement.

# Restrictions of Called USEROPEN Functions

The Intel Fortran RTL uses exactly one file descriptor per logical unit, which must be returned by the called function. Because of this, only certain system calls or library routines can be used to open the file.

System calls and library routines that do not return a file descriptor include `mknod` (see `mknod(2)`) and `fopen` (see `fopen(3)`). For example, the `fopen` routine returns a file pointer instead of a file descriptor.

# Example USEROPEN Program and Function

The following Intel Fortran code calls the USEROPEN function named UOPEN:

```
EXTERNAL   UOPEN
INTEGER    UOPEN
   .
   .
   .
OPEN (UNIT=1,FILE='ex1.dat',STATUS='NEW',USEROPEN=UOPEN,
      ERR=9,IOSTAT=errnum)
```

If the default `ifort` options are used, the external name is passed using lowercase letters with an appended trailing underscore (_). In the preceding example, the external function UOPEN would be known as `uopen_` to the linker and must be declared in C as `uopen_`.

## Compiling and Linking the C and Intel Fortran Programs

Use a single `ifort` command to compile the called `uopen_` C function `uopen_.c` and the Intel Fortran calling program `ex1.f`. The same command also links both object files by using the appropriate libraries to create the file `a.out` file, as follows:

**`ifort ex1.f uopen_.c`**

## Source Code for the C Function and Header File

The following example shows the C language function called `uopen_` and its associated header file.

```
/*
** File: uopen.h -- header file for uopen_.c
*/

#ifndef UOPEN
#define UOPEN 1

/*
**
**     Function Prototypes
**
*/
int    uopen_ (
       char  *file_name,    /* access read: name of the file
to open. */
       int   *open_flags,   /* access read: READ/WRITE, see
file.h or open(2)*/
       int   *create_mode,  /* access read: set if new file
(to be created).*/
       int   *lun,          /* access read: logical unit file
opened on.*/
       int   file_length);  /* access read: number of
characters in file_name*/

#endif

/* End of file uopen.h */


/*
** File: uopen_.c
*/

/*
** This routine opens a file using data passed by Intel
Fortran RTL.
**
**   INCLUDE FILES
*/

#include <sys/types.h>
#include <sys/stat.h>
#include <sys/file.h>
#include "uopen.h"/* Include file for this module */

int uopen_ (file_name, open_flags, create_mode, lun,
file_length)

/*
```

```
** Open a file using the parameters passed by the calling
Intel
**   Fortran program.
**
**   Formal Parameters:
*/

char  *file_name;    /* access read: name of the file to
open. */
int   *open_flags;   /* access read: READ/WRITE, see file.h
*/
int   *create_mode;  /* access read: set if new file (to be
created). */
int   *lun;          /* access read: logical unit number
file opened on. */
int   file_length;   /* access read: number of characters in
file_name. */

/*
**   Function Value/Completion Code
**
** Whatever is returned by open is immediately returned to
the
** Fortran OPEN.  The returned value is the following:
**     value >= 0 is a valid fd.
**     value <  0 is an error.
**
**   Modify open flags (logical OR) to specify the file be
opened for
**   write access only, with records appended at the end
(such as
**   writing to a shared log file).
*/

{
        int    result ;           /* Function result value */

        *open_flags =
              O_CREAT  |
              O_WRONLY |
              O_APPEND;

        result = open (file_name, *open_flags, *create_mode)
;

        return (result) ;         /* return file descriptor
or error */

        }/* End of routine uopen_ */
```

```
        /* End of file uopen_.c */
```

## Source Code for the Calling Intel Fortran Program

The following example shows the Fortran program that calls the `uopen_` C function and then performs I/O.

```
C
C  Program EX1 opens a file using USEROPEN and writes
records to it.
C  It closes and re-opens the file (without USEROPEN) and
reads 10 records.

 PROGRAM EX1

     EXTERNAL    UOPEN          ! The USEROPEN function.
     INTEGER      ERRNUM, CTR, I

 1   FORMAT (I)

     ERRNUM = 0

     WRITE (6,*) 'EX1. Access data using formatted I/O.'
     WRITE (6,*) 'EX1. Open file with USEROPEN and put some
data in it.'
     OPEN (UNIT=1, FILE='ex1.dat', STATUS='NEW',
USEROPEN=UOPEN,
          & ERR=9, & IOSTAT=errnum)

     DO CTR=1,10
       WRITE (1,1) CTR
     END DO

     WRITE (6,*) 'EX1. Close and re-open without USEROPEN.'

     CLOSE (UNIT=1)

     OPEN (UNIT=1, FILE='ex1.dat', STATUS='OLD',
          FORM='FORMATTED', & ERR=99, & IOSTAT=errnum)

     WRITE (6,*) 'EX1. Read and display what is in file.'

     DO CTR=1,10
        READ (1,1) i
        WRITE (6,*) i
     END DO

     WRITE (6,*) 'EX1.  Successful if 10 records shown.'

     CLOSE (UNIT=1,STATUS='DELETE')
```

```
     STOP

  9 WRITE (6,*) 'EX1.  Error on USEROPEN is ', errnum
     STOP

 99 WRITE (6,*) 'EX1.  Error on 2nd open is ', errnum

 END PROGRAM EX1
```

# Format of Record Types

## Fixed-Length Records

When you specify fixed-length records, all records in the file contain the same number of bytes. When you open a file that is to contain fixed-length records, you must specify the record size by using the RECL specifier. A sequentially organized file opened for direct access must contain fixed-length records, to allow the record position in the file to be computed correctly.

For relative files, the layout and overhead of fixed-length records depend on whether or not the program accessing the file was compiled with the −vms option.

For relative files where the −vms option was omitted (the default), each record has no control information.

For relative files where the −vms option was specified, each record has one byte of control information at the beginning of the record.

The figure below shows the record layout of fixed-length records:
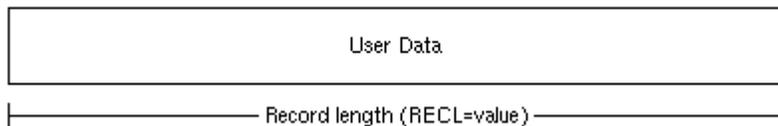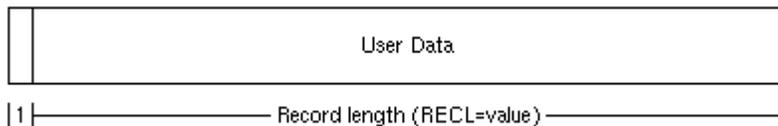


## Variable-Length Records

Variable-length records can contain any number of bytes up to a specified maximum record length, and apply only to sequential files.

Variable-length records are prefixed and suffixed by 4 bytes of control information containing length fields. The trailing length field allows a BACKSPACE request to skip back over records efficiently. The 4-byte integer value stored in each length field indicates the number of data bytes (excluding overhead bytes) in that particular variable-length record.

The character count field of a variable-length record is available when you read the record by issuing a READ statement with a Q format descriptor. You can then use the count field information to determine how many bytes should be in an I/O list.

## Variable-Length Records Less Than 2 Gigabytes

The figure below shows the record layout of variable-length records that are less than 2 gigabytes:
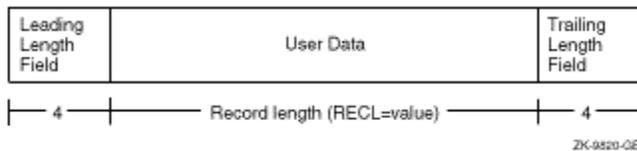


## Variable-Length Records Greater Than 2 Gigabytes

For a record length greater than 2,147,483,639 bytes, the record is divided into *subrecords*. The subrecord can be of any length from 1 to 2,147,483,639, inclusive.

The sign bit of the leading length field indicates whether the record is continued or not. The sign bit of the trailing length field indicates the presence of a preceding subrecord. The position of the sign bit is determined by the endian format of the file.

A subrecord that is continued has a leading length field with a sign bit value of 1. The last subrecord that makes up a record has a leading length field with a sign bit value of 0. A subrecord that has a preceding subrecord has a trailing length field with a sign bit value of 1. The first subrecord that makes up a record has a trailing length field with a sign bit value of 0.

The figure below shows the record layout of variable-length records that are greater than 2 gigabytes:
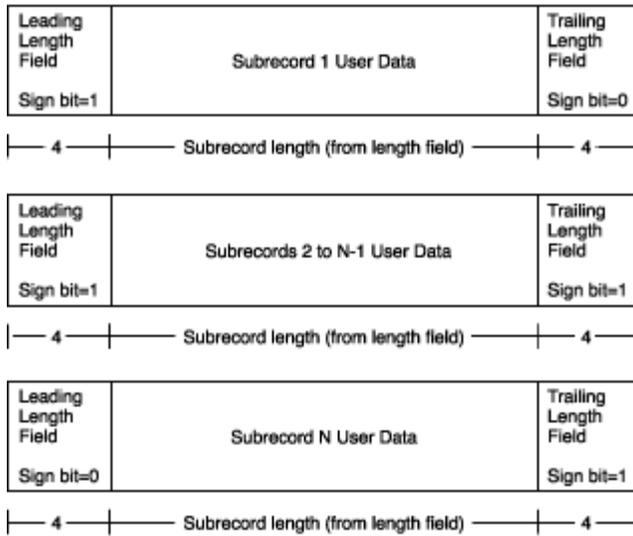
Files written with variable-length records by Intel Fortran programs usually
cannot be accessed as text files. Instead, use the Stream_LF record format for
text files with records of varying length.

# Segmented Records

A *segmented record* is a single *logical record* consisting of one or more variable-
length, unformatted records in a sequentially organized disk file. Unformatted
data written to sequentially organized files using sequential access is stored as
segmented records by default.

Segmented records are useful when you want to write exceptionally long records
but cannot or do not wish to define one long variable-length record, perhaps
because virtual memory limitations can prevent program execution. By using
smaller, segmented records, you reduce the chance of problems caused by
virtual memory limitations on systems on which the program may execute.

For disk files, the segmented record is a single logical record that consists of one
or more segments. Each segment is a *physical record*. A segmented (logical)
record can exceed the absolute maximum record length (2.14 billion bytes), but
each segment (physical record) individually cannot exceed the maximum record
length.

To access an unformatted sequential file that contains segmented records,
specify FORM= ' UNFORMATTED ' and RECORDTYPE= ' SEGMENTED ' when
you open the file.0

As shown in the figure below, the layout of segmented records consists of 4
bytes of control information followed by the user data:

ZK-9821-GE
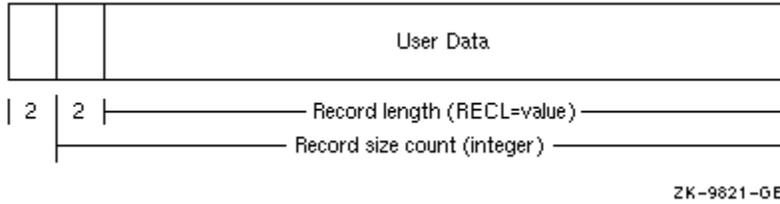
The control information consists of a 2-byte integer record size count (includes the two bytes used by the segment identifier), followed by a 2-byte integer segment identifier that identifies this segment as one of the following:

| Identifier Value | Segment Identified |
|---|---|
| 0 | One of the segments between the first and last segments |
| 1 | First segment |
| 2 | Last segment |
| 3 | Only segment |

If the specified record length is an odd number, the user data will be padded with a single blank (one byte), but this extra byte is not added to the 2-byte integer record size count.

# Stream File

A Stream file is not grouped into records and contains no control information. Stream files are used with CARRIAGECONTROL= ' NONE ' and contain character or binary data that is read or written only to the extent of the variables specified on the input or output statement.

The figure below shows the layout of a Stream file:



ZK-9822-GE
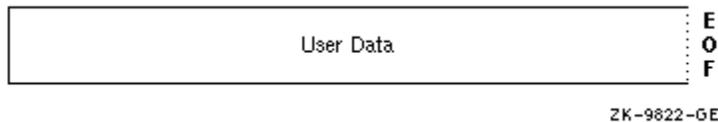
# Stream_CR and Stream_LF Records

A Stream_CR or Stream_LF record is a variable-length record whose length is indicated by explicit record terminators embedded in the data, not by a count. These terminators are automatically added when you write records to a stream-type file and are removed when you read records.

Each variety uses a different 1-byte record terminator:

- Stream_CR files use only a carriage-return as the terminator, so
  Stream_CR files must not contain embedded carriage-return characters.
- Stream_LF files use only a line-feed (new line) as the terminator, so
  Stream_LF files must not contain embedded line-feed (new line)
  characters. This is the usual operating system text file record type.

The figure below shows the layout of Stream_CR and Stream_LF records:
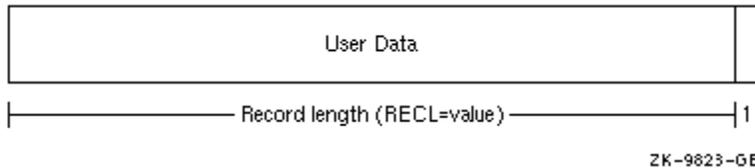


# Microsoft* Fortran PowerStation Compatible Files

When using the `-fpscomp` options for Microsoft* Fortran
PowerStation compatibility, the following types of files are
possible:

- Formatted Sequential
- Formatted Direct
- Unformatted Sequential
- Unformatted Direct

## Formatted Sequential Files

A formatted sequential file is a series of formatted records written sequentially
and read in the order in which they appear in the file. Records can vary in length
and can be empty. They are separated by carriage return (0D) and line feed (0A)
characters as shown in the following figure.

**Formatted Records in a Formatted Sequential File**
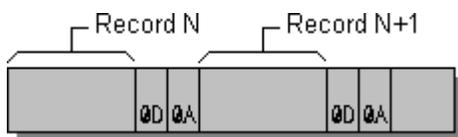


An example of a program writing three records to a formatted sequential file is
given below. The resulting file is shown in the following figure.

```
   OPEN (3, FILE='FSEQ')
! FSEQ is a formatted sequential file by default.
   WRITE (3, '(A, I3)') 'RECORD', 1
   WRITE (3, '()')
   WRITE (3, '(A11)') 'The 3rd One'
   CLOSE (3)
   END
```

**Formatted Sequential File**
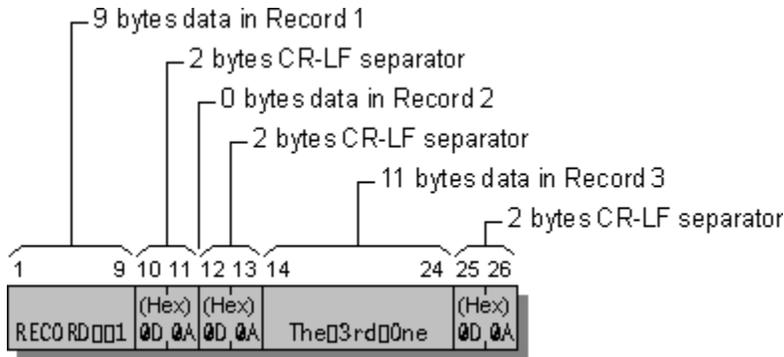


# Formatted Direct Files

In a formatted direct file, all of the records are the same length and can be written or read in any order. The record size is specified with the RECL option in an OPEN statement and should be equal to or greater than the number of bytes in the longest record.

The carriage return (CR) and line feed (LF) characters are record separators and are not included in the RECL value. Once a direct-access record has been written, you cannot delete it, but you can rewrite it.

During output to a formatted direct file, if data does not completely fill a record, the compiler pads the remaining portion of the record with blank spaces. The blanks ensure that the file contains only completely filled records, all of the same length. During input, the compiler by default also pads the input if the input list and format require more data than the record contains.

You can override the default blank padding on input by setting PAD='NO' in the OPEN statement for the file. If PAD='NO', the input record must contain the amount of data indicated by the input list and format specification. Otherwise, an error occurs. PAD='NO' has no effect on output.

An example of a program writing two records, record one and record three, to a formatted direct file is given below. The result is shown in the following figure.

```
OPEN (3,FILE='FDIR', FORM='FORMATTED', ACCESS='DIRECT',RECL=10)
WRITE (3, '(A10)', REC=1) 'RECORD ONE'
WRITE (3, '(I5)', REC=3) 30303
CLOSE (3)
END
```

**Formatted Direct File**
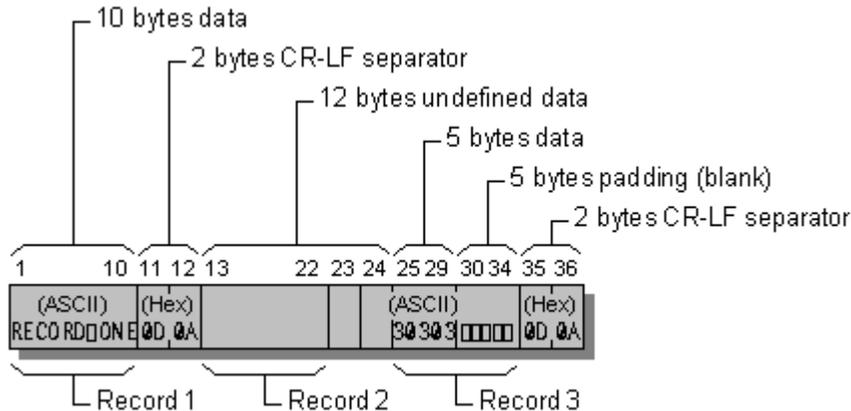


# Unformatted Sequential Files

Unformatted sequential files are organized slightly differently on different platforms. This section describes unformatted sequential files created by Intel Fortran when the `-fpscomp` option (such as `-fpscomp ioformat`) was specified. If you are accessing files from another platform that organizes them differently, see Converting Unformatted Data Overview.

The records in an unformatted sequential file can vary in length. Unformatted sequential files are organized in chunks of 130 bytes or less called physical blocks. Each physical block consists of the data you send to the file (up to 128 bytes) plus two 1-byte "length bytes" inserted by the compiler. The length bytes indicate where each record begins and ends.

A logical record refers to an unformatted record that contains one or more physical blocks. (See the following figure.) Logical records can be as big as you want; the compiler will use as many physical blocks as necessary.

When you create a logical record consisting of more than one physical block, the compiler sets the length byte to 129 to indicate that the data in the current physical block continues on into the next physical block. For example, if you write 140 bytes of data, the logical record has the structure shown in the following figure.

**Logical Record in Unformatted Sequential File**
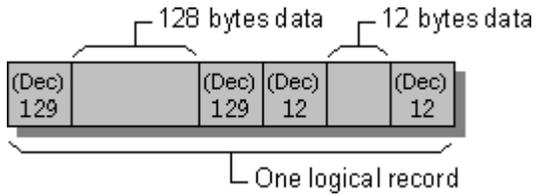


The first and last bytes in an unformatted sequential file are reserved; the first contains a value of 75, and the last holds a value of 130. Fortran uses these bytes for error checking and end-of-file references.

The following program creates the unformatted sequential file shown in the following figure:

```
!   Note: The file is sequential by default
!         -1 is FF FF FF FF hexadecimal.
!
    CHARACTER xyz(3)
    INTEGER(4) idata(35)
    DATA      idata /35 * -1/, xyz /'x', 'y', 'z'/
!
! Open the file and write out a 140-byte record:
! 128 bytes (block) + 12 bytes = 140 for IDATA, then 3 bytes for XYZ.
    OPEN (3, FILE='UFSEQ',FORM='UNFORMATTED')
    WRITE (3) idata
    WRITE (3) xyz
    CLOSE (3)
    END
```

**Unformatted Sequential File**
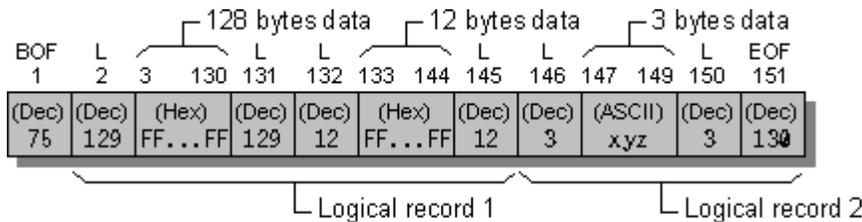


BOF   Beginning-of-file byte (75 decimal)
L     Physical-block-length byte (0 <= L <= 129)
EOF   End-of-file byte (130 decimal)

# Unformatted Direct Files

An unformatted direct file is a series of unformatted records. You can write or read the records in any order you choose. All records have the same length,

given by the RECL specifier in an OPEN statement. No delimiting bytes separate records or otherwise indicate record structure.

You can write a partial record to an unformatted direct file. Intel Visual Fortran pads these records to the fixed record length with ASCII NULL characters. Unwritten records in the file contain undefined data.

The following program creates the sample unformatted direct file shown in the following figure:

```
OPEN (3, FILE='UFDIR', RECL=10,&
  & FORM = 'UNFORMATTED', ACCESS = 'DIRECT')
WRITE (3, REC=3) .TRUE., 'abcdef'
WRITE (3, REC=1) 2049
CLOSE (3)
END
```

**Unformatted Direct File**
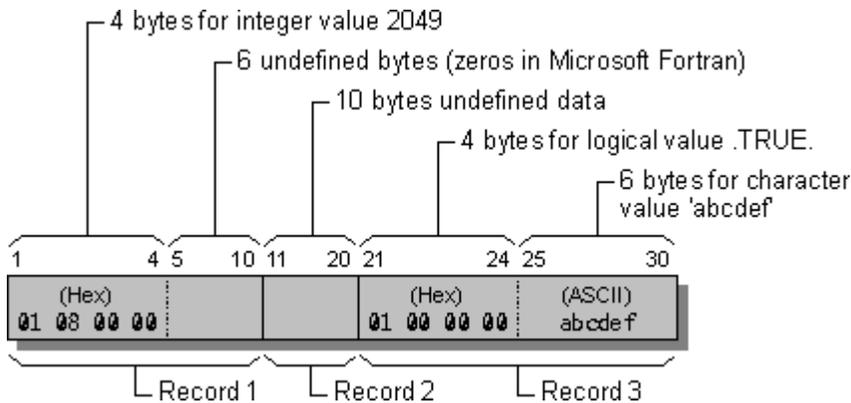


# Programming with Mixed Languages

# Programming with Mixed Languages Overview

Mixed-language programming is the process of building programs in which the source code is written in two or more languages. It allows you to:

- Call existing code that is written in another language
- Use procedures that may be difficult to implement in a particular language
- Gain advantages in processing speeds

Intel® Fortran Compiler for Linux* Systems User's Guide Volume I: Building Applications

Mixed-language programming is possible between Intel® Fortran and Intel® C++.

To properly create mixed-language programs, rules must be established for naming variables and procedures, for stack use, and for argument passing among routines written in different languages. These rules, as a whole, are the *calling convention*.

A calling convention includes:

- Stack consideration: Does a routine receive a varying or fixed number of arguments?
- Naming conventions
    - Is lowercase or uppercase significant or not significant?
    - Are external names altered?
- Argument passing protocol
    - Are arguments passed by value or by reference?
    - What are the equivalent data types and data structures among languages?

This section provides information on the calling conventions available when writing routines written in Fortran, C, C++, and assembly language. See these topics:

Calling Subprograms from the Main Program

Summary of Mixed-Language Issues

Adjusting Calling Conventions in Mixed-Language Programming Overview

Adjusting Naming Conventions in Mixed-Language Programming Overview

Protoyping a Procedure in Fortran

Exchanging and Accessing Data in Mixed-Language Programming Overview

Handling Data Types in Mixed-Language Programming Overview

Intel Fortran/C Mixed-Language Programs Overview

# Calling Subprograms from the Main Program

## Calls from the Main Program

The Intel Fortran main program can call Intel Fortran subprograms, including subprograms in static and shared libraries.

For mixed-language applications, the Intel Fortran main program can call subprograms written in Intel® C++ if the appropriate calling conventions are used (see Calling C Procedures from a Fortran program).

Intel Fortran subprograms can be called by Intel C++ main programs

## Calls to the Subprogram

You can use subprograms in static libraries if the main program is written in Intel Fortran or Intel C++.

You can use subprograms in shared libraries in mixed-language applications if the main program is written in Intel Fortran or Intel C++.

# Summary of Mixed-Language Issues

Mixed-language programming involves a call from a routine written in one language to a function, procedure, or subroutine written in another language. For example, a Fortran main program may need to execute a specific task that you want to program separately in an assembly-language procedure, or you may need to call an existing shared library or system procedure.

Mixed-language programming is possible with Intel® Fortran and Intel® C++ because each language implements functions, subroutines, and procedures in approximately the same way. The following table shows how different kinds of routines from each language correspond to each other. For example, a C main program could call an external void function, which is actually implemented as a Fortran subroutine:

**Language Equivalents for Calls to Routines**

| Language | Call with Return Value | Call with No Return Value |
| --- | --- | --- |
| Fortran | FUNCTION | SUBROUTINE |
| C and C++ | function | (void) function |

There are some important differences in the way languages implement routines. Argument passing, naming conventions, and other interface issues must be thoughtfully and consistently reconciled between any two languages to prevent

program failure and indeterminate results. However, the advantages of mixed-language programming often make the extra effort worthwhile.

A summary of a few mixed-language advantages and restrictions follows:

- Fortran/Assembly Language

  Assembly-language routines are small and execute very quickly because they do not require initialization as do high-level languages like Fortran and C. Also, they allow access to hardware instructions unavailable to the high-level language user. In a Fortran/assembly-language program, compiling the main routine in Fortran gives the assembly code access to Fortran high-level procedures and library functions, yet allows freedom to tune the assembly-language routines for maximum speed and efficiency. The main program can also be an assembly-language program.

- Fortran/C (or C++)

  Generally, Fortran/C programs are mixed to allow one to use existing code written in the other language. Either Fortran or C can call the other, so the main routine can be in either language.

This section provides an explanation of the keywords, attributes, and techniques you can use to reconcile differences between Fortran and other languages. Adjusting calling conventions, adjusting naming conventions and writing interface procedures are discussed in the next sections:

- Adjusting Calling Conventions in Mixed-Language Programming
- Adjusting Naming Conventions in Mixed-Language Programming
- Prototyping a Procedure in Fortran

After establishing a consistent interface between mixed-language procedures, you then need to reconcile any differences in the treatment of individual data types (strings, arrays, and so on). This is discussed in Exchanging and Accessing Data in Mixed-Language Programming.

### Note

This section uses the term "routine" in a generic way, to refer to functions, subroutines, and procedures from different languages.

## Adjusting Calling Conventions in Mixed-Language Programming

# Adjusting Calling Conventions in Mixed-Language Programming Overview

The calling convention determines how a program makes a call to a routine, how the arguments are passed, and how the routines are named. See Adjusting Naming Conventions in Mixed-Language Programming.

In a single-language program, calling conventions are nearly always correct, because there is one default for all routines and because header files or Fortran module files with interface blocks enforce consistency between the caller and the called routine.

In a mixed-language program, different languages cannot share the same header files. If, as a result, you link Fortran and C routines that use different calling conventions, the error is not apparent until the bad call is made at run time. During execution, the bad call causes indeterminate results and/or a fatal error, often somewhere in the program that has no apparent relation to the actual cause: memory/stack corruption due to calling errors. Therefore, you should check carefully the calling conventions for each mixed-language call.

The discussion of calling conventions between languages applies only to external procedures. You cannot call internal procedures from outside the program unit that contains them.

A calling convention affects programming in four ways:

1. The caller routine uses a calling convention to determine the order in which to pass arguments to another routine; the called routine uses a calling convention to determine the order in which to receive the arguments passed to it. In Fortran, you can specify these conventions in a mixed-language interface with the INTERFACE statement or in a data or function declaration. C/C++ and Fortran both pass arguments in order from left to right.
2. The caller routine and the called routine use a calling convention to select the option of passing a variable number of arguments.
3. The caller routine and the called routine use a calling convention to pass arguments by value (values passed) or by reference (addresses passed). Individual Fortran arguments can also be designated with ATTRIBUTES option VALUE or REFERENCE.
4. The caller routine and the called routine use a calling convention to establish naming conventions for procedure names. You can establish any procedure name you want, regardless of its Fortran name, with the ALIAS directive (or

ATTRIBUTES option ALIAS). This is useful because C is case-sensitive, while Fortran is not.

See these topics:

ATTRIBUTES Properties and Calling Conventions

Fortran/C Calling Conventions

# ATTRIBUTES Properties and Calling Conventions

The ATTRIBUTES properties (also known as options) C, REFERENCE, VALUE, and VARYING all affect the calling convention of routines. You can specify the:

- C, REFERENCE, and VARYING properties for an entire routine
- VALUE and REFERENCE properties for individual arguments

By default, Fortran passes all data by reference (except the hidden length argument of strings, which is passed by value). If the C property is used, the default changes to passing almost all data except arrays by value. However, in addition to the calling-convention property C, you can specify argument properties VALUE and REFERENCE (to pass arguments by value or by reference), regardless of the calling convention property. Arrays can only be passed by reference.

Different Fortran calling conventions can be specified by declaring the Fortran procedure to have certain attributes. Assume this example:

```
INTERFACE
   SUBROUTINE MY_SUB (I)
       !DEC$ ATTRIBUTES C, ALIAS:'My_Sub_' :: MY_SUB ! ia32 systems
       INTEGER I
   END SUBROUTINE MY_SUB
END INTERFACE
```

This code declares a subroutine named `MY_SUB` with the C property and the external name `My_Sub_` set with the ALIAS property.

For another example, the following declaration assumes the subroutine is called with the C calling convention:

```
SUBROUTINE CALLED_FROM_C (A)
```

```
!DEC$ ATTRIBUTES C :: CALLED_FROM_C
INTEGER A
```

The following table summarizes the effect of the most common Fortran calling-
convention directives:

**Calling Conventions for ATTRIBUTES Properties**

| Argument | Default | C | C, REFERENCE |
|---|---|---|---|
| Scalar | Reference | Value | Reference |
| Scalar [value] | Value | Value | Value |
| Scalar [reference] | Reference | Reference | Reference |
| String | Reference, either Len:End or Len:Mixed | String(1:1) | Reference, either Len:End or Len:Mixed |
| String [value] | Error | String(1:1) | String(1:1) |
| String [reference] | Reference, either No Len or Len:Mixed | Reference, No Len | Reference, No Len |
| Array | Reference | Reference | Reference |
| Array [value] | Error | Error | Error |
| Array [reference] | Reference | Reference | Reference |
| Derived Type | Reference | Value, size dependent | Reference |
| Derived Type [value] | Value, size dependent | Value, size dependent | Value, size dependent |
| Derived Type [reference] | Reference | Reference | Reference |
| F90 Pointer | Descriptor | Descriptor | Descriptor |
| F90 Pointer [value] | Error | Error | Error |
| F90 Pointer [reference] | Descriptor | Descriptor | Descriptor |

The procedure name is all lowercase for all the calling conventions.

The terms in the above table mean the following:

[value]          Argument assigned the VALUE attribute.

[reference]      Argument assigned the REFERENCE attribute.

Value            The argument value is pushed on the stack. All values are padded to the next 4-byte boundary.

Reference        On IA-32 systems, the 4-byte argument address is pushed on the stack.
                 On Itanium® -based systems, the 8-byte argument address is pushed on the stack.

Len:End or       For certain string arguments:
Len:Mixed

                 • Len:End applies when -nomixed_str_len_arg is set. The length of the string is pushed (by value) on the stack after all of the other arguments. This is the default.
                 • Len:Mixed applies when -mixed_str_len_arg is set. The length of the string is pushed (by value) on the stack immediately after the address of the beginning of the string.

No Len or        For certain string arguments:
Len:Mixed

                 • No Len applies when -nomixed_str_len_arg is set. The length of the string is not available to the called procedure. This is the default.
                 • Len:Mixed applies when -mixed_str_len_arg is set. The length of the string is pushed (by value) on the stack immediately after the address of the beginning of the string.

No Len           For string arguments, the length of the string is not available to the called procedure.

String(1:1)      For string arguments, the first character is converted to INTEGER(4) as in ICHAR(string(1:1)) and pushed on the stack by value.

Error            Produces a compiler error.

Descriptor       On IA-32 systems, the 4-byte address of the array descriptor. On Itanium-based systems, the 8-byte address of the array descriptor.

Size             On IA-32 systems, derived-type arguments specified by value
dependent        are passed as follows:

- Arguments from 1 to 4 bytes are passed by value.
- Arguments from 5 to 8 bytes are passed by value in two registers (two arguments).
- Arguments more than 8 bytes provide value semantics by passing a temporary storage address by reference.

The following table shows another Fortran ATTRIBUTES property that matches another language calling convention:

| Other Language Calling Convention | Matching ATTRIBUTES Property |
| --- | --- |
| C/C++ cdecl (default) | C |

The ALIAS property can be used with any other Fortran calling-convention property to preserve mixed-case names. You can also use the DECORATE property in combination with the ALIAS property to specify that the external name specified in ALIAS should have the correct prefix and postfix decorations for the calling mechanism in effect.

**Adjusting Naming Conventions in Mixed-Language Programming**

# Adjusting Naming Conventions in Mixed-Language Programming Overview

The ATTRIBUTES option C determines naming conventions as well as calling conventions.

Calling conventions specify how arguments are moved and stored; naming conventions specify how symbol names are altered when placed in a `.o` file. Names are an issue for external data symbols shared among parts of the same program as well as among external routines. Symbol names (such as the name of a subroutine) identify a memory location that must be consistent among all calling routines.

Parameter names (names given in a procedure definition to variables that are passed to it) are never affected.

Names are altered because of case sensitivity (in C), lack of case sensitivity (in Fortran), name decoration (in C++), or other issues. If naming conventions are

not reconciled, the program cannot successfully link and you will receive an "unresolved external" error.

See these topics:

C/C++ Naming Conventions

Naming Conventions for Fortran, C, and C++

Reconciling the Case of Names

Fortran Module Names and ATTRIBUTES

# C/C++ Naming Conventions

C and C++ preserve case sensitivity in their symbol tables while Fortran by default does not, a difference that requires attention. Fortunately, you can use the Fortran directive ATTRIBUTES ALIAS option to resolve discrepancies between names, to preserve mixed-case names, or to override the automatic conversion of names to all lowercase by Fortran.

C++ uses the same calling convention and argument-passing techniques as C, but naming conventions differ because of C++ decoration of external symbols. When the C++ code resides in a `.cpp` file (created when you select C/C++ file from the integrated development environment), C++ name decoration semantics are applied to external names, often resulting in linker errors. The extern "C" syntax makes it possible for a C++ module to share data and routines with other languages by causing C++ to drop name decoration.

The following example declares `prn` as an external function using the C naming convention. This declaration appears in C++ source code:

```
extern "C" { void prn(); }
```

To call functions written in Fortran, declare the function as you would in C and use a "C" linkage specification. For example, to call the Fortran function FACT from C++, declare it as follows:

```
extern "C" { int FACT( int n ); }
```

The extern "C" syntax can be used to adjust a call from C++ to other languages, or to change the naming convention of C++ routines called from other languages. However, extern "C" can only be used from within C++. If the C++ code does not use extern "C" and cannot be changed, you can call C++ routines only by determining the name decoration and generating it from the other language.

186

Such an approach should only be used as a last resort, because the decoration scheme is not guaranteed to remain the same between versions.

Use of extern "C" has some restrictions:

- You cannot declare a member function with extern "C".
- You can specify extern "C" for only one instance of an overloaded function; all other instances of an overloaded function have C++ linkage.

# Procedure Names in Fortran, C, and C++

The following table summarizes how Fortran, C, and C++ handle procedure names:

| Language | Attributes | Name Translated As | Case of Name in .o File |
|---|---|---|---|
| Fortran | *c*DEC$ ATTRIBUTES C | *name_* | All lowercase |
| Fortran | default | *name_* | All lowercase |
| C | `cdecl` (default) | *name_* | Mixed case preserved |
| C | `__stdcall` | *_name@n* | Mixed case preserved |
| C++ | Default | *_name@@decoration* | Mixed case preserved |

# Reconciling the Case of Names

The following summarizes how to reconcile names between languages:

- All-lowercase names

  If the name of the routine appears as all lowercase in C, then naming conventions are automatically correct. Any case can be used in the Fortran source code, including mixed case, since the name is changed to all lowercase.

- Mixed-case names

  If the name of a routine appears as mixed-case in C and you cannot change the name, then you can resolve this naming conflict by using the Fortran ATTRIBUTES ALIAS option ALIAS is required in this situation because otherwise Fortran will not preserve the mixed-case name.

  To use the ALIAS option, place the name in quotation marks exactly as it is to appear in the `.o` file.
  The following is an example for referring to the C function `My_Proc`:
  ```
  !DEC$ ATTRIBUTES ALIAS:'My_Proc_' :: My_Proc
  ```

# Fortran Module Names and ATTRIBUTES

Fortran module entities (data and procedures) have external names that differ from other external entities. Module names use the convention:

*MODULENAME*`_mp_`*ENTITY_*

*MODULENAME* is the name of the module and *ENTITY* is the name of the module procedure or module data contained within *MODULENAME*. `_mp_` is the separator between the module and entity names and is always lowercase.

For example:

```
MODULE mymod
   INTEGER a
CONTAINS
   SUBROUTINE b (j)
      INTEGER j
   END SUBROUTINE
END MODULE
```

This results in the following symbols being defined in the compiled `.o` file:

```
mymod_mp_a_
mymod_mp_b_
```

Compiler options can affect the naming of module data and procedures.

### 📝Note

Except for ALIAS, ATTRIBUTES properties do not affect the module name.

The following table shows how each ATTRIBUTES property affects the
subroutine in the previous example module.

**Effect of ATTRIBUTES Options on Fortran Module Names**

| ATTRIBUTES Property Given to Routine 'b' | Procedure Name in .o file |
| --- | --- |
| None | `mymod_mp_b_` |
| C | `mymod_mp_b_` |
| ALIAS | Overrides all others, name as given in the alias |
| VARYING | No effect on name |

You can write code to call Fortran modules or access module data from other
languages. As with other naming and calling conventions, the module name must
match between the two languages. Generally, this means using the C convention
in Fortran, and if defining a module in another language, using the ALIAS
property to match the name within Fortran. For examples, see Using Modules in
Fortran/C Mixed-Language Programming.

# Prototyping a Procedure in Fortran

You define a prototype (interface block) in your Fortran source code to tell the
Fortran compiler which language conventions you want to use for an external
reference. The interface block is introduced by the INTERFACE statement. See
"Program Units and Procedures" in the *Language Reference* for a description of
the INTERFACE statement.

The general form for the INTERFACE statement is:

```
INTERFACE
routine statement
[routine ATTRIBUTE options]
[argument ATTRIBUTE options]
formal argument declarations
END routine name
END INTERFACE
```

The `routine statement` defines either a FUNCTION or a SUBROUTINE,
where the choice depends on whether a value is returned or not, respectively.
The optional `routine ATTRIBUTE options` (such as C) determine the calling,

naming, and argument-passing conventions for the routine in the prototype statement. The optional *argument ATTRIBUTE options* (such as VALUE and REFERENCE) are properties attached to individual arguments. The *formal argument declarations* are Fortran data type declarations. Note that the same INTERFACE block can specify more than one procedure.

For example, suppose you are calling a C function that has the following prototype:

```
extern void My_Proc (int i);
```

The Fortran call to this function should be declared with the following INTERFACE block:

```
INTERFACE
   SUBROUTINE my_Proc (I)
     !DEC$ ATTRIBUTES C, ALIAS:'My_Proc_' :: my_Proc
     INTEGER I
   END SUBROUTINE my_Proc
END INTERFACE
```

Note that, except in the ALIAS string, the case of `My_Proc` in the Fortran program does not matter.

**Exchanging and Accessing Data in Mixed-Language Programming**

# Exchanging and Accessing Data in Mixed-Language Programming Overview

You can use several approaches to sharing data between mixed-language routines, which can be used within the individual languages as well.

Generally, if you have a large number of parameters to work with or you have a large variety of parameter types, you should consider using modules or external data declarations. This is true when using any given language, and to an even greater extent when using mixed languages.

See also Using Modules in Fortran/C Mixed-Language Programming.

See these topics:

Passing Arguments in Mixed-Language Programming

Using Common External Data in Mixed-Language Programming

# Passing Arguments in Mixed-Language Programming

You can pass data between Fortran, C, and C++ through calling argument lists just as you can within each language (for example, the argument list a, b and c in CALL MYSUB(a,b,c)). There are two ways to pass individual arguments:

- *By value*, which passes the argument's value.
- *By reference*, which passes the address of the arguments. On IA-32 systems, Fortran, C, and C++ use 4-byte addresses. On Itanium® -based systems, these languages use 8-byte addresses.

You need to make sure that for every call, the calling program and the called routine agree on how each argument is passed. Otherwise, the called routine receives bad data.

The Fortran technique for passing arguments changes depending on the calling convention specified. By default, Fortran passes all data by reference (except the hidden length argument of strings, which is passed by value).

If the ATTRIBUTES C option is used, the default changes to passing all data by value except arrays. If the procedure has the REFERENCE option as well as the C option, all arguments by default are passed by reference.

In Fortran, in addition to establishing argument passing with the calling-convention option C, you can specify argument options, VALUE and REFERENCE, to pass arguments by value or by reference. In mixed-language programming, it is a good idea to specify the passing technique explicitly rather than relying on defaults.

### Note

In addition to ATTRIBUTES, the `-[no]mixed_str_len_arg` compiler option also establishes some default argument passing conventions (such as for hidden length of strings).

Examples of passing by reference and value for C follow. All are interfaces to the example Fortran subroutine TESTPROC below. The definition of TESTPROC declares how each argument is passed. The REFERENCE option is not strictly necessary

in this example, but using it makes the argument's passing convention conspicuous.

```
SUBROUTINE TESTPROC( VALPARM, REFPARM )
    !DEC$ ATTRIBUTES VALUE :: VALPARM
    !DEC$ ATTRIBUTES REFERENCE :: REFPARM
    INTEGER VALPARM
    INTEGER REFPARM
END SUBROUTINE
```

In C and C++ all arguments are passed by value, except arrays, which are passed by reference to the address of the first member of the array. Unlike Fortran, C and C++ do not have calling-convention directives to affect the way individual arguments are passed. To pass non-array C data by reference, you must pass a pointer to it. To pass a C array by value, you must declare it as a member of a structure and pass the structure. The following C declaration sets up a call to the example Fortran `testproc` subroutine:

```
extern void testproc( int ValParm, int *RefParm );
```

The following table summarizes how to pass arguments by reference and value. An array name in C is equated to its starting address because arrays are normally passed by reference. You can assign the REFERENCE property to a procedure, as well as to individual arguments.

**Passing Arguments by Reference and Value**

| Language | ATTRIBUTE | Argument Type | To Pass by Reference | To Pass by Value |
|---|---|---|---|---|
| Fortran | Default | Scalars and derived types | Default | VALUE option |
| | C option | Scalars and derived types | REFERENCE option | Default |
| | Default | Arrays | Default | Cannot pass by value |
| | C option | Arrays | Default | Cannot pass by value |
| C/C++ | | Non-arrays | Pointer argument_name | Default |
| | | Arrays | Default | Struct {type} array_name |

This table does not describe argument passing of strings and Fortran 95/90 pointer arguments in Intel Fortran, which are constructed differently than other arguments. By default, Fortran passes strings by reference along with the string length. String length placement depends on whether the compiler option `-mixed_str_len_arg` (immediately after the address of the beginning of the string) or `-nomixed_str_len_arg` (after all arguments) is set. The default setting is `-nomixed_str_len_arg`.

Fortran 95/90 array pointers and assumed-shape arrays are passed by passing the address of the array descriptor.

For a discussion of the effect of attributes on passing Fortran 95/90 pointers and strings, see Handling Fortran 90 Pointers and Allocatable Arrays and Handling Character Strings.

# Using Common External Data in Mixed-Language Programming

Common external data structures include Fortran common blocks, and C structures and variables that have been declared global or external. All of these data specifications create external variables, which are variables available to routines outside the routine that defines them.

External variables are case-sensitive, so the cases must be matched between different languages, as discussed in the section on naming conventions. Common external data exchange is described in the following sections:

- Using Global Variables
- Using Fortran Common Blocks and C Structures

## Using Global Variables in Mixed-Language Programming

A variable can be shared between Fortran and C by declaring it as global (or COMMON) in one language and accessing it as an external variable in the other language. In Fortran programs, variables must be passed as arguments.

In Fortran, a variable can access a global parameter by using the EXTERN option for ATTRIBUTES. For example:

```
!DEC$ ATTRIBUTES C, EXTERN :: idata
INTEGER idata (20)
```

EXTERN tells the compiler that the variable is actually defined and declared global in another source file. If Fortran declares a variable external with EXTERN, the language it shares the variable with must declare the variable global.

In C, a variable is declared global with the statement:

```
int idata[20]; // declared as global (outside of any function)
```

Fortran can declare the variable global (COMMON) and other languages can reference it as external:

```
! Fortran declaring PI global
REAL PI
COMMON /PI/ PI ! Common Block and variable have the same name
```

In C, the variable is referenced as an external with the statement:

```
//C code with external reference to PI0
extern float PI;
```

Note that the global name C references is the name of the Fortran common block, not the name of a variable within a common block. Thus, you cannot use blank common to make data accessible between C and Fortran. In the preceding example, the common block and the variable have the same name, which helps keep track of the variable between the two languages. Obviously, if a common block contains more than one variable they cannot all have the common block name. (See Using Fortran Common Blocks and C Structures.)

## Using Fortran Common Blocks and C Structures

To reference C structures from Fortran common blocks and vice versa, you must take into account the way the common blocks and structures differ in their methods of storing member variables in memory. Fortran places common block variables into memory in order as close together as possible, with the following rules:

- A single BYTE, INTEGER(1), LOGICAL(1), or CHARACTER variable in common block list begins immediately following the previous variable or array in memory.

- All other types of single variables begin at the next even address immediately following the previous variable or array in memory.
- All arrays of variables begin on the next even address immediately following the previous variable or array in memory, except for CHARACTER arrays which always follow immediately after the previous variable or array.
- All common blocks begin on a four-byte aligned address.

Because of these padding rules, you must consider the alignment of C structure elements with Fortran common block elements and assure matching either by making all variables the same types and kinds in both languages (using only 4-byte and 8-byte data types in both languages simplifies this) or by using the C pack pragmas in the C code around the C structure to make C data packing like Fortran's. For example:

```
#pragma pack(2)
struct {
        int N;
        char INFO[30];
} examp;
#pragma pack()
```

To restore the original packing, you must add `#pragma pack( )` at the end of the structure. (Remember: Fortran module data can be shared directly with C structures with appropriate naming.)

Once you have dealt with alignment and padding, you can give C access to an entire common block or set of common blocks. Alternatively, you can pass individual members of a Fortran common block in an argument list, just as you can any other data item. Use of common blocks for mixed-language data exchange is discussed in the following sections:

- Accessing Common Blocks and C Structures Directly
- Passing the Address of a Common Block

## Accessing Common Blocks and C Structures Directly

You can access Fortran common blocks directly from C by defining an external C structure with the appropriate fields, and making sure that alignment and padding between Fortran and C are compatible. The C and ALIAS ATTRIBUTES options can be used with a common block to allow mixed-case names.

As an example, suppose your Fortran code has a common block named `Really`, as shown:

```
    !DEC$ ATTRIBUTES ALIAS:'Really' :: Really
    REAL(4) x, y, z(6)
    REAL(8) ydbl
    COMMON / Really / x, y, z(6), ydbl
```

You can access this data structure from your C code with the following external data structure:

```
#pragma pack(2)
extern struct {
  float x, y, z[6];
  double ydbl;
} Really;
#pragma pack()
```

You can also access C structures from Fortran by creating common blocks that correspond to those structures. This is the reverse case from that just described. However, the implementation is the same because after common blocks and structures have been defined and given a common address (name), and assuming the alignment in memory has been dealt with, both languages share the same memory locations for the variables.

## Passing the Address of a Common Block

To pass the address of a common block, simply pass the address of the first variable in the block, that is, pass the first variable by reference. The receiving C or C++ module should expect to receive a structure by reference.

In the following example, the C function `initcb` receives the address of a common block with the first variable named `n`, which it considers to be a pointer to a structure with three fields:

**Fortran source code**:

```
!
  INTERFACE
    SUBROUTINE initcb (BLOCK)
      !DEC$ ATTRIBUTES C :: initcb
      !DEC$ ATTRIBUTES REFERENCE :: BLOCK
      INTEGER BLOCK
    END SUBROUTINE
  END INTERFACE
!
  INTEGER n
  REAL(8) x, y
  COMMON /CBLOCK/n, x, y
    .  .  .
  CALL initcb( n )
```

**C source code**:

```
//
#pragma pack(2)
struct block_type
{
  int n;
  double x;
  double y;
};
#pragma pack()
//
void initcb( struct block_type *block_hed )
 {
   block_hed->n = 1;
   block_hed->x = 10.0;
   block_hed->y = 20.0;
}
```

**Handling Data Types in Mixed-Language Programming**

# Handling Data Types in Mixed-Language Programming Overview

Even when you have reconciled calling conventions, naming conventions, and methods of data exchange, you must still be concerned with data types, because each language handles them differently.

The following table lists the equivalent data types between Fortran and C:

**Equivalent Data Types**

| Fortran Data Type | C Data Type |
|---|---|
| INTEGER(1) | `char` |
| INTEGER(2) | `short` |
| INTEGER(4) | `int, long` |
| INTEGER(8) | `_int64` |
| REAL(4) | `float` |
| REAL(8) | `double` |

| REAL(16) | --- |
|----------|-----|
| CHARACTER(1) | `unsigned char` |
| CHARACTER*(*) | See Handling Character Strings |
| COMPLEX(4) | `struct complex4`<br>`{`<br>`float real,`<br>`imag;`<br>`};` |
| COMPLEX(8) | `struct complex8`<br>`{`<br>`double real,`<br>`imag;`<br>`};` |
| COMPLEX(16) | --- |
| All LOGICAL types | Use integer types for C |

See these topics:

Handling Numeric, Complex, and Logical Data Types

Handling Fortran Array Pointers and Allocatable Arrays

Handling Intel Fortran Pointers

Handling Arrays and Fortran Array Descriptors

Handling Character Strings

Handling User-Defined Types

# Handling Numeric, Complex, and Logical Data Types

Normally, passing numeric data does not present a problem. If a C program passes an unsigned data type to a Fortran routine, the routine can accept the argument as the equivalent signed data type, but you should be careful that the range of the signed type is not exceeded.

The table of Equivalent Data Types summarizes equivalent numeric data types for Fortran and C/C++.

C and C++ do not directly implement the Fortran types COMPLEX(4), COMPLEX(8), and COMPLEX(16). However, you can write structures that are equivalent. The type COMPLEX(4) has two fields, both of which are 4-byte floating-point numbers; the first contains the real-number component, and the second contains the imaginary-number component. The type COMPLEX is equivalent to the type COMPLEX(4). The types COMPLEX(8) and COMPLEX(16) are similar except that each field contains an 8-byte or 16-byte floating-point number respectively.

## Note

On IA-32 systems, Fortran functions of type COMPLEX place a hidden COMPLEX argument at the beginning of the argument list. C functions that implement such a call from Fortran must declare this hidden argument explicitly, and use it to return a value. The C return type should be void.

Following are the C/C++ structure definitions for the Fortran COMPLEX types:

```
struct complex4 {
   float real, imag;
};
struct complex8 {
   double real, imag;
};
```

A Fortran LOGICAL(2) is stored as a 2-byte indicator value (0=false, and the –fpscomp [no]logicals compiler option determines how true values are handled). A Fortran LOGICAL(4) is stored as a 4-byte indicator value, and LOGICAL(1) is stored as a single byte. The type LOGICAL is the same as LOGICAL(4), which is equivalent to type int in C.

You can use a variable of type LOGICAL in an argument list, module, common block, or global variable in Fortran and type int in C for the same argument. Type LOGICAL(4) is recommended instead of the shorter variants for use in common blocks.

The Intel C++ class type has the same layout as the corresponding C struct type, unless the class defines virtual functions or has base classes. Classes that lack those features can be passed in the same way as C structures.

# Returning Complex Type Data

If a Fortran program expects a procedure to return a COMPLEX DOUBLE COMPLEX value, the Fortran compiler adds an additional argument to the beginning of the called procedure argument list. This additional argument is a pointer to the location where the called procedure must store its result.

Example below shows the Fortran code for returning a complex data type procedure called `WBAT` and the corresponding C routine.

**Example of Returning Complex Data Types from C to Fortran**

**Fortran code:**

```
COMPLEX BAT, WBAT
REAL X, Y
BAT = WBAT ( X, Y )
```

**Corresponding C routine:**

```
struct _mycomplex { float real, imag };
typedef struct _mycomplex _single_complex;

void WBAT (_single_complex location, float *x, float *y)
{
float realpart;
float imaginarypart;
... program text, producing realpart and imaginarypart...
*location.real = realpart;
*location.imag = imaginarypart;
}
```

In the above example, the following restrictions and behaviors apply:

- The argument location does not appear in the Fortran call; it is added by the compiler.
- The C subroutine must copy the result's real and imaginary parts correctly into `location`.
- The called procedure is type `void`.

If the function returned a DOUBLE COMPLEX value, the type `float` would be replaced by the type `double` in the definition of `location` in `WBAT`.

# Handling Fortran Array Pointers and Allocatable Arrays

How Fortran 95/90 array pointers and arrays are passed is affected by the ATTRIBUTES properties in effect, and by the INTERFACE, if any, of the procedure they are passed to.

If the INTERFACE declares the array pointer or array with deferred shape (for example, ARRAY(:)), its descriptor is passed. This is true for array pointers and all arrays, not just allocatable arrays. If the INTERFACE declares the array pointer or array with fixed shape, or if there is no interface, the array pointer or array is passed by base address as a contiguous array, which is like passing the first element of an array for contiguous array slices.

When a Fortran 95/90 array pointer or array is passed to another language, either its descriptor or its base address can be passed.

The following shows how allocatable arrays and Fortran 95/90 array pointers are passed with different attributes in effect:

- If the property of the array pointer or array is not included or is REFERENCE, it is passed by descriptor, regardless of the property of the passing procedure (None; C; or C, REFERENCE).
- If the property of the array pointer or array is VALUE, an error is returned, regardless of the property of the passing procedure.

Note that the VALUE option cannot be used with descriptor-based arrays.

When you pass a Fortran array pointer or an array by descriptor to a non-Fortran routine, that routine needs to know how to interpret the descriptor. Part of the descriptor is a pointer to address space, as a C pointer, and part of it is a description of the pointer or array properties, such as its rank, stride, and bounds.

For information about the Intel Fortran array descriptor format, see Handling Arrays and Fortran Array Descriptors.

Fortran 95/90 pointers that point to scalar data contain the address of the data and are not passed by descriptor.

# Handling Integer Pointers

*Integer pointers* (also known as Cray*-style pointers) are not the same as Fortran 90 pointers, but are instead like C pointers. Integer pointers are 4-byte INTEGER quantities on IA-32 systems, and 8-byte INTEGER quantities on Itanium®-based systems.

## Passing Integer Pointers

When passing an integer pointer to a routine written in another language:

- The argument should be declared in the non-Fortran routine as a pointer
  of the appropriate data type.
- The argument passed from the Fortran routine should be the integer
  pointer name, not the pointee name.

For example:

```
! Fortran main program.
      INTERFACE
         SUBROUTINE Ptr_Sub (p)
         !DEC$ ATTRIBUTES C, ALIAS:'Ptr_Sub' :: Ptr_Sub
            INTEGER p
         END SUBROUTINE Ptr_Sub
      END INTERFACE
      REAL A(10), VAR(10)
      POINTER (p, VAR) ! VAR is the pointee
                       ! p is the integer pointer
      p = LOC(A)
      CALL Ptr_Sub (p)
      WRITE(*,*) 'A(4) = ', A(4)
      END
  !
  //C subprogram
      void Ptr_Sub (float *p)
      {
         p[3] = 23.5;
      }
```

On Itanium-based systems, the declaration for `p` in the INTERFACE block should
be `INTEGER(8) p`.

When the main Fortran program and C function are built and executed, the
following output appears:

```
A(4) = 23.50000
```

# Receiving Pointers

When receiving a pointer from a routine written in another language:

- The argument should be declared in the non-Fortran routine as a pointer
  of the appropriate data type and passed as usual.
- The argument received by the Fortran routine should be declared as an
  integer pointer name, and the POINTER statement should associate it
  with a pointee of the appropriate data type (matching the data type of the

passing routine). When inside the Fortran routine, use the pointee to set and access what the pointer points to.

For example:

```
! Fortran subroutine.
      SUBROUTINE Iptr_Sub (p)
      !DEC$ ATTRIBUTES C, ALIAS:'Iptr_Sub' :: Iptr_Sub
         integer VAR(10)
         POINTER (p, VAR)
         OPEN (8, FILE='STAT.DAT')
         READ (8, *) VAR(4) ! Read from file and store the
                            ! fourth element of VAR
      END SUBROUTINE Iptr_Sub
!
//C main program
extern void Iptr_Sub(int *p);
main ( void )
{
  int a[10];
  Iptr_Sub (&a[0]);
  printf("a[3] = %i\n", a[3]);
}
```

When the main C program and Fortran subroutine are built and executed, the following output appears if the STAT.DAT file contains 4:

```
  a[3] = 4
```

# Handling Arrays and Fortran Array Descriptors

Fortran 95/90 allows arrays to be passed as array elements, as array subsections, or as whole arrays referenced by array name. Within Fortran, array elements are ordered in column-major order, meaning the subscripts of the lowest dimensions vary first.

When using arrays between Fortran and another language, differences in element indexing and ordering must be taken into account. You must reference the array elements individually and keep track of them. Fortran and C vary in the way that array elements are indexed. Array indexing is a source-level consideration and involves no difference in the underlying data.

Fortran and C arrays differ in two ways:

- The value of the lower array bound is different. By default, Fortran indexes the first element of an array as 1. C and C++ index it as 0. Fortran subscripts should therefore be one higher. (Fortran also provides the option of specifying another integer lower bound.)
- In arrays of more than one dimension, Fortran varies the left-most index the fastest, while C varies the right-most index the fastest. These are sometimes called column-major order and row-major order, respectively.

In C, the first four elements of an array declared as X[3][3] are:

```
X[0][0] X[0][1] X[0][2] X[1][0]
```

In Fortran, the first four elements are:

```
X(1,1) X(2,1) X(3,1) X(1,2)
```

The order of indexing extends to any number of dimensions you declare. For example, the C declaration:

```
int arr1[2][10][15][20];
```

is equivalent to the Fortran declaration:

```
INTEGER arr1( 20, 15, 10, 2 )
```

The constants used in a C array declaration represent extents, not upper bounds as they do in other languages. Therefore, the last element in the C array declared as int arr[5][5] is arr[4][4], not arr[5][5].

The following table shows equivalencies for array declarations.

**Equivalent Array Declarations for Different Languages**

| Language | Array Declaration | Array Reference from Fortran |
|---|---|---|
| Fortran | DIMENSION x(i, k) -or- *type* x(i, k) | x(i, k) |
| C/C++ | *type* x[ k ] [ i ] | x( i -1, k -1) |

# Intel Fortran Array Descriptor Format

For cases where Fortran 95/90 needs to keep track of more than a pointer memory address, the Intel Fortran Compiler uses an *array descriptor*, which stores the details of how an array is organized.

When using an explicit interface (by association or procedure interface block), Intel Fortran generates a descriptor for the following types of array arguments:

- Pointers to arrays (array pointers)
- Assumed-shape arrays

Certain data structure arguments do not use a descriptor, even when an appropriate explicit interface is provided. For example, explicit-shape and assumed-size arrays do not use a descriptor. In contrast, array pointers and allocatable arrays use descriptors regardless of whether they are used as arguments.

When calling between Intel Fortran and a non-Fortran language (such as C), using an *implicit interface* allows the array argument to be passed without an Intel Fortran descriptor. However, for cases where the called routine needs the information in the Intel Fortran descriptor, declare the routine with an *explicit interface* and specify the dummy array as either an assumed-shape array or with the pointer attribute.

You can associate a Fortran 95/90 pointer with any piece of memory, organized in any way desired (so long as it is "rectangular" in terms of array bounds). You can also pass Fortran 95/90 pointers to other languages, such as C, and have the other language correctly interpret the descriptor to obtain the information it needs.

However, using array descriptors can increase the opportunity for errors and the corresponding code is not portable. In particular, be aware of the following:

- If the descriptor is not defined correctly, the program may access the wrong memory address, possibly causing a General Protection Fault.
- Array descriptor formats are specific to each Fortran compiler. Code that uses array descriptors is not portable to other compilers or platforms. For example, the current Intel Fortran array descriptor format differs from the array descriptor format for Intel Fortran 7.0.
- The array descriptor format may change in the future.

The components of the current Intel Fortran array descriptor on IA-32 systems are as follows:

- The first longword (bytes 0 to 3) contains the base address. The base address plus the offset defines the first memory location (start) of the array.
- The second longword (bytes 4 to 7) contains the size of a single element of the array.
- The third longword (bytes 8 to 11) contains the offset. The offset is added to the base address to define the start of the array.
- The fourth longword (bytes 12 to 15) contains the low-order bit set if the array has been defined (storage allocated). Other bits may also be set by the compiler within this longword, for example, to indicate a contiguous array.
- The fifth longword (bytes 16 to 19) contains the number of dimensions (rank) of the array.
- The sixth longword (bytes 20 to 23) is reserved.
- The remaining longwords (bytes 24 up to 107) contain information about each dimension (up to seven). Each dimension is described by three additional longwords:
    - The number of elements (extent)
    - The distance between the starting address of two successive elements in this dimension, in bytes.
    - The lower bound

An array of rank one requires three additional longwords for a total of nine longwords (6 + 3*1) and ends at byte 35. An array of rank seven is described in a total of 27 longwords (6 + 3*7) and ends at byte 107.

For example, consider the following declaration:

```
integer,target :: a(10,10)
```

```
integer,pointer :: p(:,:)
```

```
p => a(9:1:-2,1:9:3)
```

```
call f(p)
```

.

.

.

The descriptor for actual argument p would contain the following values:

- The first longword (bytes 0 to 3) contain the base address (assigned at run-time).
- The second longword (bytes 4 to 7) is set to 4 (size of a single element).
- The third longword (bytes 8 to 11) contain the offset (assigned at run-time).
- The fourth longword (bytes 12 to 15) contains 1 (low bit is set).
- The fifth longword (bytes 16 to 19) contains 2 (rank).
- The sixth longword is reserved.
- The seventh, eighth, and ninth longwords (bytes 24 to 35) contain information for the first dimension, as follows:
  - 5 (extent)
  - -8 (distance between elements)
  - 9 (the lower bound)
- For the second dimension, the tenth, eleventh, and twelfth longwords (bytes 36 to 47) contain:
  - 3 (extent)
  - 120 (distance between elements)
  - 1 (the lower bound)
- Byte 47 is the last byte for this example.

### Note

The format for the descriptor on Itanium-based systems is identical to that on IA-32 systems, except that all fields are 8-bytes long, instead of 4-bytes.

# Handling Character Strings

By default, Intel Fortran passes a hidden length argument for strings. The hidden length argument consists of an unsigned 4-byte integer (IA-32 systems) or unsigned 8-byte integer (Itanium® -based systems), always passed by value, added to the end of the argument list. You can alter the default way strings are passed by using attributes.

The following table shows the effect of various attributes on passed strings:

**Effect of ATTRIBUTES Properties on Character Strings Passed as Arguments**

| Argument | Default | C | C, REFERENCE |
|---|---|---|---|
| String | Passed by | First character | Passed by reference. |

| | reference, along with length | converted to INTEGER(4) and passed by value | along with length |
|---|---|---|---|
| String with VALUE option | Error | First character converted to INTEGER(4) and passed by value | First character converted to INTEGER(4) and passed by value |
| String with REFERENCE option | Passed by reference, possibly along with length | Passed by reference, no length | Passed by reference, no length |

The important things to note about the above table are:

- Character strings without the VALUE or REFERENCE attribute that are passed to C routines are not passed by reference. Instead, only the first character is passed and it is passed by value.
- Character strings with the VALUE option passed to C routines are not passed by reference. Instead, only the value of the first character is passed.
- For string arguments with default ATTRIBUTES, ATTRIBUTES C, or REFERENCE:
    - When `-nomixed_str_len_arg` is set, the length of the string is pushed (by value) on the stack after all of the other arguments. This is the default.
    - When `-mixed_str_len_arg` is set, the length of the string is pushed (by value) on the stack immediately after the address of the beginning of the string.
- For string arguments passed by reference with default ATTRIBUTES:
    - When `-nomixed_str_len_arg` is set, the length of the string is not available to the called procedure. This is the default.
    - When `-mixed_str_len_arg` is set, the length of the string is pushed (by value) on the stack immediately after the address of the beginning of the string.

Since all strings in C are pointers, C expects strings to be passed by reference, without a string length. In addition, C strings are null-terminated while Fortran strings are not. There are two basic ways to pass strings between Fortran and C: convert Fortran strings to C strings, or write C routines to accept Fortran strings.

To convert a Fortran string to C, choose a combination of attributes that passes the string by reference without length, and null terminate your strings. For example:

```
INTERFACE
```

```
     SUBROUTINE Pass_Str (string)
       !DEC$ ATTRIBUTES C, DECORATE,ALIAS:'Pass_Str' :: Pass_Str
       CHARACTER*(*) string
       !DEC$ ATTRIBUTES REFERENCE :: string
     END SUBROUTINE
   END INTERFACE
   CHARACTER(40) forstring
   DATA forstring /'This is a null-terminated string.'C/
```

The following example shows the extension of using the null-terminator for the string in
the Fortran DATA statement (see "C Strings" in the *Intel® Fortran Language
Reference*):

```
  DATA forstring /'This is a null-terminated string.'C/
```

The C interface is:

```
  void Pass_Str (char *string)
```

To get your C routines to accept Fortran strings, C must account for the length argument
passed along with the string address. For example:

```
  ! Fortran code
  INTERFACE
     SUBROUTINE Pass_Str (string)
     CHARACTER*(*) string
  END INTERFACE
```

The C routine must expect two arguments:

```
  void pass_str (char *string, unsigned int length_arg )
```

This interface handles the hidden-length argument, but you must still reconcile C
strings that are null-terminated and Fortran strings that are not. In addition, if the
data assigned to the Fortran string is less than the declared length, the Fortran
string will be blank padded.

Rather than trying to handle these string differences in your C routines, the best
approach in Fortran/C mixed programming is to adopt C string behavior
whenever possible.

Fortran functions that return a character string using the syntax CHARACTER*(*)
place a hidden string argument and the length of the string at the beginning of
the argument list.

C functions that implement such a Fortran function call must declare this hidden
string argument explicitly and use it to return a value. The C return type should
be void. However, you are more likely to avoid errors by not using character-

string return functions. Use subroutines or place the strings into modules or global variables whenever possible.

# Returning Character Data Types

If a Fortran program expects a function to return data of type `CHARACTER`, the Fortran compiler adds two additional arguments to the beginning of the called procedure's argument list:

- The first argument is a pointer to the location where the called procedure should store the result.
- The second is the maximum number of characters that must be returned, padded with white spaces if necessary.

The called routine must copy its result through the address specified in the first argument. Example that follows shows the Fortran code for a return character function called `MAKECHARS` and corresponding C routine.

**Example of Returning Character Types from C to Fortran**

```
Fortran code
CHARACTER*10 CHARS, MAKECHARS
DOUBLE PRECISION X, Y
CHARS = MAKECHARS( X, Y )
Corresponding C Routine
void makechars ( result, length, x, y );
char *result;
int length;
double *x, *y;
{
...program text, producing returnvalue...
for (i = 0; i < length; i++ ) {
result[i] = returnvalue[i];
}
}
```

In the above example, the following restrictions and behaviors apply:

- The function's length and result do not appear in the call statement; they are added by the compiler.
- The called routine must copy the result string into the location specified by `result`; it must not copy more than `length` characters.
- If fewer than `length` characters are returned, the return location should be padded on the right with blanks; Fortran does not use zeros to terminate strings.
- The called procedure is type `void`.

- You must use lowercase names for C routines and `INTERFACE` blocks to make the calls using lower case.

# Handling User-Defined Types

Fortran 95/90 supports *user-defined types* (data structures similar to C structures). User-defined types can be passed in modules and common blocks just as other data types, but the other language must know the type's structure.

For example:

**Fortran Code:**

```
TYPE LOTTA_DATA
    SEQUENCE
    REAL A
    INTEGER B
    CHARACTER(30) INFO
    COMPLEX CX
    CHARACTER(80) MOREINFO
 END TYPE LOTTA_DATA
 TYPE (LOTTA_DATA) D1, D2
 COMMON /T_BLOCK/ D1, D2
```

In the Fortran code above, the SEQUENCE statement preserves the storage order of the derived-type definition.

**C Code:**

```
/* C code accessing D1 and D2 */
extern struct {
    struct {
        float a;
        int b;
        char info[30];
        struct {
            float real, imag;
            } cx;
        char moreinfo[80];
    } d1, d2;
} t_block;
```

## Intel Fortran/C Mixed-Language Programs

# Intel Fortran/C Mixed-Language Programs Overview

See these topics:

Compiling and Linking Mixed-Language Programs

Using Modules in Fortran/C Mixed-Language Programming

Calling C Procedures from an Intel Fortran Program

# Compiling and Linking Intel Fortran/C Programs

Your application can contain both C and Fortran source files. If your main program is a Fortran source file (`myprog.for`) that calls a routine written in C (`cfunc.c`), you can use the following sequence of commands to build your application:

```
icc -c cfunc.c
ifort -o myprog myprog.for cfunc.o
```

The `icc` (for Intel C++) command compiles `cfunc.c`. The `-c` option specifies that the linker is not called. This command creates `cfunc.o`. The `ifort` command compiles `myprog.for` and links `cfunc.o` with the object file created from `myprog.for` to create `myprog`.

# Using Modules in Fortran/C Mixed-Language Programming

Modules are the simplest way to exchange large groups of variables with C, because Intel Fortran modules are directly accessible from C/C++.

The following example declares a module in Fortran, then accesses its data from C:

**Fortran code**

```
! F90 Module definition
   MODULE EXAMP
      REAL A(3)
      INTEGER I1, I2
      CHARACTER(80) LINE
      TYPE MYDATA
         SEQUENCE
         INTEGER N
         CHARACTER(30) INFO
      END TYPE MYDATA
   END MODULE EXAMP
```

## C code

```
\* C code accessing module data *\
extern float examp_mp_a[3];
extern int examp_mp_i1, examp_mp_i2;
extern char examp_mp_line[80];
extern struct {
        int n;
        char info[30];
} examp_mp_mydata;
```

When the C++ code resides in a `.cpp` file, C++ semantics are applied to external names, often resulting in linker errors. In this case, use the extern "C" syntax (see C/C++ Naming Conventions):

```
\* C code accessing module data in .cpp file*\
extern "C" float examp_mp_a[3];
extern "C" int examp_mp_i1, examp_mp_i2;
extern "C" char examp_mp_line[80];
extern "C" struct {
        int n;
        char info[30];
} examp_mp_mydata;
```

You can also define a module procedure in C and make that routine part of a Fortran module by using the ALIAS directive. The C code is:

```
 // C procedure
void pythagoras (float a, float b, float *c)
{
    *c = (float) sqrt(a*a + b*b);
}
```

Using the same example when the C++ code resides in a `.cpp` file, use the extern "C" syntax (see C/C++ Naming Conventions):

```
// C procedure
extern "C" void pythagoras (float a, float b, float *c)
{
    *c = (float) sqrt(a*a + b*b);
}
```

The Fortran code to define the module CPROC:

```
! Fortran 95/90 Module including procedure
   MODULE CPROC
      INTERFACE
         SUBROUTINE PYTHAGORAS (a, b, res)
          !DEC$ ATTRIBUTES C :: PYTHAGORAS
          !DEC$ ATTRIBUTES REFERENCE :: res
! res is passed by REFERENCE because its individual attribute
! overrides the subroutine's C attribute
            REAL a, b, res
! a and b have the VALUE attribute by default because
! the subroutine has the C attribute
         END SUBROUTINE
      END INTERFACE
   END MODULE
```

The Fortran code to call this routine using the module CPROC:

```
! Fortran 95/90 Module including procedure
   USE CPROC
      CALL PYTHAGORAS (3.0, 4.0, X)
      TYPE *,X
   END
```

# Calling C Procedures from an Intel Fortran Program

## Naming Conventions

By default, the Fortran compiler converts function and subprogram names to upper case. The C compiler never performs case conversion. A C procedure called from a Fortran program must, therefore, be named using the appropriate case. For example, consider the following calls:

| CALL PROCNAME() | The C procedure must be named PROCNAME. |
| X=FNNAME() | The C procedure must be named FNNAME |

In the first call, any value returned by PROCNAME is ignored. In the second call to a function, FNNAME must return a value.

## Passing Arguments Between Fortran and C Procedures

By default, Fortran subprograms pass arguments by reference; that is, they pass a pointer to each actual argument rather than the value of the argument. C programs, however, pass arguments by value. Consider the following:

- When a Fortran program calls a C function, the C function's formal arguments must be declared as pointers to the appropriate data type.
- When a C program calls a Fortran subprogram, each actual argument must be specified explicitly as a pointer.

## Error Handling

# Error Handling Overview

See these topics:

Run-Time Library Default Error Processing

Run-Time Environment Variables

Handling Run-Time Errors

Signal Handling

Overriding the Default Run-Time Library Exception Handler

Obtaining Traceback Information with TRACEBACKQQ

# Run-Time Library Default Error Processing

During execution, your program may encounter errors or exception conditions. These conditions can result from any of the following:

- Errors that occur during I/O operations
- Invalid input data
- Argument errors in calls to the mathematical library
- Arithmetic errors
- Other system-detected errors

The Intel® Fortran Run-Time Library (RTL) generates appropriate messages and takes action to recover from errors whenever possible.

A default action is defined for each error recognized by the Fortran RTL. The default actions described throughout this chapter occur unless overridden by explicit error-processing methods.

The way in which the Fortran RTL actually processes errors depends upon the following factors:

- The severity of the error. For instance, the program usually continues executing when an error message with a severity level of warning or info (informational) is detected.
- For certain errors associated with I/O statements, whether or not an I/O error-handling specifier was specified.
- For certain errors, whether or not the default action of an associated signal was changed.
- For certain errors related to arithmetic operations (including floating-point exceptions), compilation options can determine whether the error is reported and the severity of the reported error.

How arithmetic exception conditions are reported and handled depends on the cause of the exception and how the program was compiled. Unless the program was compiled to handle exceptions, the exception might not be reported until after the instruction that caused the exception condition. The following compiler options are related to handling errors and exceptions:

- The `-check bounds` option generates extra code to catch certain conditions.
- The `-check noformat` and `-check nooutput_conversion` options reduce the severity level of the associated run-time error to allow program continuation.
- The `-fpen` options control the handling and reporting of floating-point arithmetic exceptions at run time.
- The `-warn xxxx`, `-u`, `-nowarn -w`, and `-wl` options control compile-time warning messages, which in some circumstances can help determine the cause of a run-time error.

# Run-Time Message Format

When errors occur during execution (run time) of a program, the Fortran RTL issues diagnostic messages. These run-time messages have the following format:

`forrtl:` *`severity`* `(`*`nnn`*`):` *`message-text`*

where:

- `forrtl` identifies the source as the Intel Fortran RTL.
- *`severity`* identifies the severity level: `severe`, `error`, `warning`, or `info`.
- *`nnn`* identifies the message number; also the IOSTAT value for I/O statements.
- *`message-text`* explains the event that caused the message.

The severity levels are described in order of greatest to least severity:

- A `severe` message must be corrected. The program's execution is terminated when the error is encountered, unless the program's I/O statements use the END, EOR, or ERR branch specifiers to transfer control, perhaps to a routine that uses the IOSTAT specifier.
- An `error` message should be corrected. The program might continue execution, but the output from this execution may be incorrect.
- A `warning` message should be investigated. The program continues execution, but output from this execution may be incorrect.
- An `info` message is for informational purposes only. The program continues.

For severe errors, stack trace information is produced by default, unless the environment variable `FOR_DISABLE_STACK_TRACE` is set. If the command-line option `-traceback` is set, the stack trace information contains program counters set to symbolic information. Otherwise, the information contains merely hexadecimal program counter information.

In some cases, stack trace information is also produced by the compiled code at run time to provide details about the creation of array temporaries.

If `FOR_DISABLE_STACK_TRACE` is set, no stack trace information is produced.

See the following example of stack trace information. The program generates an error at line 12:

```
program ovf
```

```
real*4 x(5),y(5)
integer*4 i

x(1) = -1e32
x(2) = 1e38
x(3) = 1e38
x(4) = 1e38
x(5) = -36.0

do i=1,5
   y(i) = 100.0*(x(i))
   print *, 'x = ', x(i), ' x*100.0 = ',y(i)
end do

end


> ifort -O0 -fpe0 -traceback ovf.f90 -o ovf.exe
> ovf.exe

 x =  -1.0000000E+32  x*100.0 =  -1.0000000E+34     (1)
forrtl: error (72): floating overflow
Image         PC            Routine        Line        Source
ovf.exe       08049E4A    MAIN__              14    ovf.f90
ovf.exe       08049F08    Unknown       Unknown     Unknown
ovf.exe       400B3507    Unknown       Unknown     Unknown
ovf.exe       08049C51    Unknow        Unknown     Unknown
Abort

> setenv FOR_DISABLE_STACK_TRACE true
> ovf.exe

 x =  -1.0000000E+32  x*100.0 =  -1.0000000E+34
forrtl: error (72): floating overflow              (2)
Abort
```

The following information corresponds to the numbers at the right of the example:

(1) Stack trace information when the traceback information is present.

(2) No stack trace information, because the `FOR_DISABLE_STACK_TRACE` environment variable is set.

# Message Catalog File Location

The Intel Fortran RTL uses a message catalog file to store the text associated with each run-time message. When a run-time error occurs, the Fortran RTL searches for the message catalog file at the following location (assuming that the default location was chosen during installation):

```
/usr/lib/ifcore_msg.cat
```

However, if the environment variable NLSPATH is set to point to another location, the RTL uses that location to find the message catalog file.

When using the C shell, you can set NLSPATH as follows:

**setenv NLSPATH /<installation-path>/%N**

Before executing an Intel Fortran program on a system where Intel Fortran is not installed, you need to copy the redistributable files from the appropriate locations specified in the fredist.txt file.

When a run-time error occurs on a system where the message file is not found, the following messages may appear:

```
forrtl: info: Fortran error message number is nnn.
forrtl: warning: Could not open message catalog:
ifcore_msg.cat.
forrtl: info: Check environment variable NLSPATH and
protection of usr/lib/ifcore_msg.cat
```

The Intel Fortran RTL returns an error number (displayed after the severity level) that the calling program can use with an IOSTAT variable to handle various I/O conditions.

For more information on NLSPATH, see the reference page environ(5).

# Values Returned to the Shell at Program Termination

An Intel Fortran program can terminate in one of several ways:

- The program runs to normal completion. A value of zero is returned to the shell.
- The program stops with a STOP or a PAUSE statement. A value of zero is returned to the shell.
- The program stops because of a signal that is caught but does not allow the program to continue. A value of 1 is returned to the shell.
- The program stops because of a severe run-time error. The error number for that run-time error is returned to the shell.
- The program stops with a CALL EXIT statement. The value passed to EXIT is returned to the shell.

# Forcing a Core Dump for Severe Errors

You can force a core dump for severe errors that do not usually cause a `core` file to be created. Before running the program, set the `decfort_dump_flag` environment variable to any of the common TRUE values (Y, y, Yes, yEs, True, and so forth) to cause severe errors to create a `core` file. For instance, the following C shell command sets the `decfort_dump_flag` environment variable:

**setenv decfort_dump_flag y**

The `core` file is written to the current directory and can be examined using a debugger.

### 📒 Note

If you requested a core file to be created on severe errors and you don't get one when expected, the problem might be that your process limit for the allowable size of a core file is set too low (or to zero). See the man page for your shell for information on setting process limits. For example, the C shell command `limit` (with no arguments) will report your current settings, and `limit coredumpsize unlimited` will raise the allowable limit to your current system maximum.

# Handling Run-Time Errors

Whenever possible, the Intel Fortran RTL does certain error handling, such as generating appropriate messages and taking necessary action to recover from errors. You can explicitly supplement or override default actions by using the following methods:

- To transfer control to error-handling code within the program, use the ERR, EOR, and END branch specifiers in I/O statements.
- To identify Fortran-specific I/O errors based on the value of Intel Fortran RTL error codes, use the I/O status specifier (IOSTAT) in I/O statements (or call the ERRSNS subroutine).
- Obtain system-level error codes by using the appropriate library routines.
- For certain error conditions, use the signal handling facility to change the default action to be taken.

These error-processing methods are complementary; you can use any or all of them within the same program to obtain Intel Fortran run-time and Linux* system error codes.

## Using the END, EOR, and ERR Branch Specifiers

When a severe error occurs during Intel Fortran program execution, the default action is to display an error message and terminate the program. To override this default action, there are three branch specifiers you can use in I/O statements to transfer control to a specified point in the program:

- The END branch specifier handles an end-of-file condition.
- The EOR branch specifier handles an end-of-record condition for nonadvancing reads.
- The ERR branch specifier handles all error conditions.

If you use the END, EOR, or ERR branch specifiers, no error message is displayed and execution continues at the designated statement, usually an error-handling routine.

You might encounter an unexpected error that the error-handling routine cannot handle. In this case, do one of the following:

- Modify the error-handling routine to display the error message number
- Remove the END, EOR, or ERR branch specifiers from the I/O statement that causes the error

After you modify the source code, compile, link, and run the program to display the error message. For example:

```
READ (8,50,ERR=400)
```

If any severe error occurs during execution of this statement, the Intel Fortran RTL transfers control to the statement at label 400. Similarly, you can use the END specifier to handle an end-of-file condition that might otherwise be treated as an error. For example:

```
READ (12,70,END=550)
```

When using nonadvancing I/O, use the EOR specifier to handle the end-of-record condition. For example:

```
150 FORMAT (F10.2, F10.2, I6)
  READ (UNIT=20, FMT=150, SIZE=X, ADVANCE='NO', EOR=700) A,
F, I
```

You can also use ERR as a specifier in an OPEN, CLOSE, or INQUIRE statement. For example:

```
OPEN (UNIT=10, FILE='FILNAM', STATUS='OLD', ERR=999)
```

If an error is detected during execution of this OPEN statement, control transfers to the statement at label 999.

# Using the IOSTAT Specifier

You can use the IOSTAT specifier to continue program execution after an I/O error and to return information about I/O operations. Certain errors are not returned in IOSTAT.

The IOSTAT specifier can supplement or replace the END, EOR, and ERR branch transfers. Execution of an I/O statement containing the IOSTAT specifier suppresses the display of an error message and defines the specified integer variable, array element, or scalar field reference as one of the following:

- A value of -2 if an end-of-record condition occurs with nonadvancing reads.
- A value of -1 if an end-of-file condition occurs.
- A value of 0 for normal completion (not an error condition, end-of-file, or end-of-record condition).
- A positive integer value if an error condition occurs. (This value is one of the Fortran-specific IOSTAT numbers listed in the run-time error message. See un-Time Error Messages.

Following the execution of the I/O statement and assignment of an IOSTAT value, control transfers to the END, EOR, or ERR statement label, if any. If there is no control transfer, normal execution continues.

You can include `/opt/intel_fc_80/include/for_iosdef.for` in your program to obtain symbolic definitions for the values of IOSTAT.

The following example uses the IOSTAT specifier and the `for_iosdef.for` file to handle an OPEN statement error (in the FILE specifier):

```fortran
      CHARACTER(LEN=40) :: FILNM
      INCLUDE 'for_iosdef.for'

      DO I=1,4
        FILNM = ''
        WRITE (6,*)  'Type file name '
        READ (5,*) FILNM
        OPEN (UNIT=1, FILE=FILNM, STATUS='OLD', IOSTAT=IERR,
ERR=100)
        WRITE (6,*) 'Opening file: ', FILNM
!        (process the input file)
        CLOSE (UNIT=1)
        STOP
```

```
  100  IF (IERR .EQ. FOR$IOS_FILNOTFOU) THEN
          WRITE (6,*) 'File: ', FILNM, ' does not exist '
        ELSE IF (IERR .EQ. FOR$IOS_FILNAMSPE) THEN
           WRITE (6,*) 'File: ', FILNM, ' was bad, enter new
file name'
        ELSE
          PRINT *, 'Unrecoverable error, code =', IERR
          STOP
        END IF
      END DO
      WRITE (6,*) 'File not found. Type ls to find file and
run again'
   END PROGRAM
```

Another way to obtain information about an error is the ERRSNS subroutine,
which allows you to obtain the last I/O system error code associated with an Intel
Fortran RTL error (see the *Intel Fortran Language Reference*).

# Signal Handling

A *signal* is an abnormal event generated by one of various sources, such as:

- A user of a terminal
- Program or hardware error
- Request of another program
- When a process is stopped to allow access to the control terminal

You can optionally set certain events to issue signals, for example:

- When a process resumes after being stopped
- When the status of a child process changes
- When input is ready at the terminal

Some signals terminate the receiving process if no action is taken (optionally
creating a `core` file), while others are simply ignored unless the process has
requested otherwise.

Except for certain signals, calling the `signal` or `sigaction` routine allows
specified signals to be ignored or causes an interrupt (transfer of control) to the
location of a user-written signal handler.

You can establish one of the following actions for a signal with a call to `signal`:

- Ignore the specified signal (identified by number).

- Use the default action for the specified signal, which can reset a previously established action.
- Transfer control from the specified signal to a procedure to receive the signal, specified by name.

Calling the `signal` routine lets you change the action for a signal, such as intercepting an operating system signal and preventing the process from being stopped.

The table below shows the signals that the Intel Fortran RTL arranges to catch when a program is started:

| Signal | Intel Fortran RTL message |
|--------|---------------------------|
| SIGFPE | `Floating-point exception` (number 75) |
| SIGINT | `Process interrupted` (number 69) |
| SIGIOT | `IOT trap signal` (number 76) |
| SIGQUIT | `Process quit` (number 79) |
| SIGSEGV | `Segmentation fault` (number 174) |
| SIGTERM | `Process killed` (number 78) |

Calling the `signal` routine (specifying the numbers for these signals) results in overwriting the signal-handling facility set up by the Intel Fortran RTL. The only way to restore the default action is to save the returned value from the first call to `signal`.

When using a debugger, it may be necessary to enter a command to allow the Intel Fortran RTL to receive and handle the appropriate signals.

# Overriding the Default Run-Time Library Exception Handler

To override the default run-time library exception handler, your application must call `signal` to change the action for the signal of interest.

For example, assume that you want to change the signal action to cause your application to call `abort()` and generate a `core` file.

The following example adds a function named `clear_signal_` to call `signal()` and change the action for the SIGABRT signal:

```
#include <signal.h>
void clear_signal_()
```

```
{
        signal (SIGABRT, SIG_DFL);
}
int myabort_()
{
        abort();
        return 0;
}
```

A call to the `clear_signal()` local routine must be added to `main`. Make sure
that the call appears before any call to the local `myabort()` routine:

```
program aborts
integer i

call clear_signal()

i = 3
if (i < 5) then
     call myabort()
end if
end
```

# Obtaining Traceback Information with TRACEBACKQQ

You can obtain traceback information in your application by calling the
TRACEBACKQQ routine.

TRACEBACKQQ allows an application to initiate a stack trace. You can use this
routine to report application detected errors, use it for debugging, and so on. It
uses the standard stack trace support in the Intel Fortran run-time system to
produce the same output that the run-time system produces for unhandled errors
and exceptions (severe error message). The TRACEBACKQQ subroutine
generates a stack trace showing the program call stack as it was leading up to
the point of the call to TRACEBACKQQ.

The error message string normally included from the run-time support is replaced
with the user-supplied message text or omitted if no user string is specified.
Traceback output is directed to the target destination appropriate for the
application type, just as it is when traceback is initiated internally by the run-time
support.

In the most simple case, a user can generate a stack trace by coding the call to
TRACEBACKQQ with no arguments:

```
CALL TRACEBACKQQ()
```

This call causes the run-time library to generate a traceback report with no leading header message, from wherever the call site is, and terminate execution.

You can specify arguments that generate a stack trace with the user-supplied string as the header and instead of terminating execution, return control to the caller to continue execution of the application. For example:

```
CALL TRACEBACKQQ(STRING="Done with pass 1",USER_EXIT_CODE=-1)
```

By specifying a user exit code of -1, control returns to the calling program. Specifying a user exit code with a positive value requests that specified value be returned to the operating system. The default value is 0, which causes the application to abort execution.

## Using Libraries

# Using Libraries Overview

See these topics:

Libraries Provided by Intel Fortran

Portability Library Overview

Math Libraries Overview

# Libraries Provided by Intel Fortran

Intel Fortran provides different types of libraries, such as static or DLL, single-threaded or multi-threaded, for certain libraries.

The table below shows the libraries provided by the compiler:

| File | Description |
|------|-------------|
| `crtxi.o` | |
| `crtxn.o` | |
| `for_main.o` | |
| `icrt.internal.map` | |
| `icrt.link` | |
| `ifcore_msg.cat` | |
| `libcprts.a` | C++ standard language library. |

| | |
|---|---|
| `libcprts.so`<br>`libcprts.so.5` | |
| `libcxa.a`<br>`libcxa.so`<br>`libcxa.so.5` | C++ language library indicating I/O data location. |
| `libcxaguard.a`<br>`libcxaguard.so`<br>`libcxaguard.so.5` | |
| `libguide.a`<br>`libguide.so` | OpenMP* static library for the parallelizer tool |
| `libguide_stats.a`<br>`libguide_stats.so` | Support for parallelizer tool with performance and profile information |
| `libifcore.a`<br>`libifcore.so`<br>`libifcore.so.5` | Intel-specific Fortran run-time library |
| `libifcoremt.a`<br>`libifcoremt.so`<br>`libifcoremt.so.5` | Multithreaded Intel-specific Fortran run-time library |
| `libifport.a`<br>`libifport.so`<br>`libifport.so.5` | Portability and POSIX support |
| `libimf.a`<br>`libimf.so` | Math library |
| `libirc.a` | Intel-specific library (optimizations) |
| `libircmt.a` | Multithreaded Intel-specific library (optimizations) |
| `libompstub.a` | Library that resolves references to OMP subroutines when OMP is not in use |
| `libsvml.a` | Short vector math library |
| `libunwind.a`<br>`libunwind.so`<br>`libunwind.so.5` | Unwind support |

**Portability Library**

# Portability Library Overview

Intel® Fortran includes functions and subroutines that ease porting of code to or from a PC, or allow you to write code on a PC that is compatible with other platforms. The portability library is called `libifport.a`. Frequently used functions are included in a portability module called IFPORT.

See these topics:

Using the Portability Library libifport.a

# Using the Portability Library libifport.a

You can use the portability library `libifport.a` in one of two ways:

- Add the statement USE IFPORT to your program. This statement includes the portability library `libifport.a`.
- Call portability routines using the correct parameters and return value.

`libifport.a` is passed to the linker by default during linking. To prevent `libifport.a` from being passed to the linker, specify the `-fpscomp nolibs` option.

Using the `libifport.a` portability library provides interface blocks and parameter definitions for the routines, as well as compiler verification of calls.

Some routines in this library can be called with different sets of arguments, and sometimes even as a function instead of a subroutine. In these cases, the arguments and calling mechanism determine the meaning of the routine. The `libifport.a` portability library contains generic interface blocks that give procedure definitions for these routines.

Fortran 95/90 contains intrinsic procedures for many of the portability functions. The portability routines are extensions to the Fortran 95 standard. When writing new code, use Fortran 95/90 intrinsic procedures whenever possible (for portability and performance reasons).

# Portability Routines

This section describes some of the portability routines and how to use them.

For a complete list of the routines, see the table of Portability Routines in the Overview chapter of the *Intel Fortran Libraries Reference*.

## Information Retrieval Routines

Information retrieval routines return information about system commands, command-line arguments, environment variables, and process or user information.

Group, user, and process ID are INTEGER(4) variables. Login name and host name are character variables. The functions GETGID and GETUID are provided for portability, but always return 1.

# Process Control Routines

Process control routines control the operation of a process or subprocess. You can wait for a subprocess to complete with either SLEEP or ALARM, monitor its progress and send signals via KILL, and stop its execution with ABORT.

In spite of its name, KILL does not necessarily stop execution of a program. Rather, the routine signaled could include a handler routine that examines the signal and takes appropriate action depending on the code passed.

Note that when you use SYSTEM, commands are run in a separate shell. Defaults set with the SYSTEM function, such as current working directory or environment variables, do not affect the environment the calling program runs in.

The portability library does not include the FORK routine. On Linux* systems, FORK creates a duplicate image of the parent process. Child and parent processes each have their own copies of resources, and become independent from one another.

# Numeric Values and Conversion Routines

Numeric values and conversion routines are available for calculating Bessel functions, data type conversion, and generating random numbers.

Some of these functions have equivalents in standard Fortran 95/90. Data object conversion can be accomplished by using the INT intrinsic function instead of LONG or SHORT. The intrinsic subroutines RANDOM_NUMBER and RANDOM_SEED perform the same functions as the random number functions listed in the table showing numeric values and conversion routines.

Other bit manipulation functions such as AND, XOR, OR, LSHIFT, and RSHIFT are intrinsic functions. You do not need the IFPORT module to access them. Standard Fortran 95/90 includes many bit operation routines. These routines are listed in Chapter 9 of the *Language Reference*, in table 9-2, under Category Bit.

# Input and Output Routines

The portability library contains routines that change file properties, read and write characters and buffers, and change the offset position in a file. These input and output routines can be used with standard Fortran input or output statements

such as READ or WRITE on the same files, provided that you take into account the following:

- When used with direct files, after an FSEEK, GETC, or PUTC operation, the record number is the number of the next whole record. Any subsequent normal Fortran I/O to that unit occurs at the next whole record. For example, if you seek to absolute location 1 of a file whose record length is 10, the NEXTREC returned by an INQUIRE would be 2. If you seek to absolute location 10, NEXTREC would still return 2.
- On units with CARRIAGECONTROL='FORTRAN' (the default), PUTC and FPUTC characters are treated as carriage control characters if they appear in column 1.
- On sequentially formatted units, the C string "\n"c, which represents the carriage return/line feed escape sequence, is written as CHAR(13) (carriage return) and CHAR(10) (line feed), instead of just line feed, or CHAR(10). On input, the sequence 13 followed by 10 is returned as just 10. (The length of character string "\n"c is 1 character, whose ASCII value, indicated by ICHAR('\n'c), is 10.)
- Reading and writing is in a raw form for direct files. Separators between records can be read and overwritten. Therefore, be careful if you continue using the file as a direct file.

I/O errors arising from the use of these routines result in an Intel Visual Fortran run-time error.

Some portability file I/O routines have equivalents in standard Fortran 95/90. For example, you could use the ACCESS function to check a file specified by name for accessibility according to mode. It tests a file for read, write, or execute permission, as well as checking to see if the file exists. It works on the file attributes as they exist on disk, not as a program's OPEN statement specifies them.

Instead of ACCESS, you can use the INQUIRE statement with the ACTION parameter to check for similar information. (The ACCESS function always returns 0 for read permission on FAT files, meaning that all files have read permission.)

# Date and Time Routines

Various date and time routines are available to determine system time, or convert it to local time, Greenwich Mean Time, arrays of date and time elements, or an ASCII character string.

DATE and TIME are available as either a function or subroutine. Because of the name duplication, if your programs do not include the USE IFPORT statement, each separately compiled program unit can use only one of these versions. For

example, if a program calls the subroutine TIME once, it cannot also use TIME
as a function.

Standard Fortran 95/90 includes date and time intrinsic subroutines.

# Error Handling Routines

Error handling routines detect and report errors.

IERRNO error codes are analogous to `errno` on UNIX systems. The IFPORT
module provides parameter definitions for many of UNIX's `errno` names, found
typically in `errno.h` on UNIX systems.

IERRNO is updated only when an error occurs. For example, if a call to the
GETC function results in an error, but two subsequent calls to PUTC succeed, a
call to IERRNO returns the error for the GETC call. Examine IERRNO
immediately after returning from one of the portability library routines. Other
standard Fortran 90 routines might also change the value to an undefined value.

If your application uses multithreading, remember that IERRNO is set on a per-
thread basis.

# System, Drive, or Directory Control and Inquiry Routines

You can retrieve information about devices, directories, and files with these
routines.

Standard Fortran 90 provides the INQUIRE statement, which returns detailed file
information either by file name or unit number. Use INQUIRE as an equivalent to
FSTAT, LSTAT or STAT. LSTAT and STAT return the same information; STAT is
the preferred function.

# Additional Routines

You can also use portability routines for program call and control, keyboards and
speakers, file management, arrays, floating-point inquiry and control, IEEE*
functionality, and other miscellaneous uses. See the table of Portability Routines
in the Overview chapter of the *Intel Fortran Libraries Reference*.

# Math Libraries

`libimf.a` is the math library provided by Intel and `libm.a` is the math library provided with gcc*.

Both of these libraries are linked in by default on IA-32 and Itanium® -based compilers. Both libraries are linked in because there are math functions supported by the GNU math library that are not in the Intel math library. This linking arrangement allows the GNU users to have all functions available when using `ifort`, with Intel optimized versions available when supported.

`libimf.a` is linked in before `libm.a`. If you link in `libm.a` first, it will change the versions of the math functions that are used.

It is recommended that you place `libimf.a` in the first directory specified in the `LD_LIBRARY_PATH` variable. The `libimf.a` and `libm.a` libraries are always linked with Fortran programs.

For example, if you place a library in directory `/perform/`, set the `LD_LIBRARY_PATH` variable to specify a list of directories, containing all other libraries, separated by semicolons.

### libimf.a on the IA-32 Compiler

For the IA-32 compiler, `libimf.a` contains both generic math routines and versions of the math routines optimized for special use with the Intel Pentium® 4 and Intel® Xeon™ processors.

### libimf.a on the Itanium-Based Compiler

For the Itanium-based compiler, `libimf.a` is optimized for use with the Itanium architecture. The compiler provides inlined versions of math library primitives and schedules the generated code with surrounding instructions. This can improve the performance of typical floating-point applications.

## Reference Information

# Compile-Time Environment Variables

The compile-time environment variables are:

- `FPATH`
  The path for include and module files.
- `IFORTCFG`
  The configuration file to use instead of the default configuration file.

- `LD_LIBRARY_PATH`
  The path for shared (.so) library files.
- `PATH`
  The path for compiler executable files.
- `TMP, TMPDIR, TEMP`
  Specifies the directory in which to store temporary files. See Temporary Files Created by the Compiler or Linker.

# Run-Time Environment Variables

The Intel Fortran run-time system recognizes several environment variables. These variables can be used to customize run-time diagnostic error reporting, for example.

The run-time environment variables are:

- `decfort_dump_flag`
  If this variable is set to Y or y, a core dump will be taken when any severe Intel Fortran run-time error occurs.
- `F_UFMTENDIAN`
  This variable specifies the numbers of the units to be used for little-endian-to-big-endian conversion purposes. See Environment Variable F_UFMTENDIAN Method.
- `FOR_ACCEPT`
  The ACCEPT statement does not include an explicit logical unit number. Instead, it uses an implicit internal logical unit number and the `FOR_ACCEPT` environment variable. If `FOR_ACCEPT` is not defined, the code `ACCEPT f,iolist` reads from `stdin` (standard input). If `FOR_ACCEPT` is defined (as a file name optionally containing a path), the specified file would be read.
- `FOR_DIAGNOSTIC_LOG_FILE`
  If this variable is set to the name of a file, diagnostic output is written to the specified file.
  The Fortran run-time system attempts to open that file (append output) and write the error information (ASCII text) to the file.
  The setting of `FOR_DIAGNOSTIC_LOG_FILE` is independent of `FOR_DISABLE_DIAGNOSTIC_DISPLAY`, so you can disable the screen display of information but still capture the error information in a file. The text string you assign for the file name is used literally, so you must specify the full name. If the file open fails, no error is reported and the run-time system continues diagnostic processing.
- `FOR_DISABLE_DIAGNOSTIC_DISPLAY`
  Disables the display of all error information.
  This variable is helpful if you just want to test the error status of your

program and do not want the Fortran run-time system to display any information about an abnormal program termination.

- `FOR_DISABLE_STACK_TRACE`
  This variable disables the call stack trace information that follows the displayed severe error message text.
  The Fortran run-time error message is displayed whether or not `FOR_DISABLE_STACK_TRACE` is set to true.
- `FOR_IGNORE_EXCEPTIONS`
  This variable disables the default run-time exception handling, for example, to allow just-in-time debugging. The run-time system exception handler returns EXCEPTION_CONTINUE_SEARCH to the operating system, which looks for other handlers to service the exception.
- `FOR_NOERROR_DIALOGS`
  This variable disables the display of dialog boxes when certain exceptions or errors occur. This is useful when running many test programs in batch mode to prevent a failure from stopping execution of the entire test stream.
- `FOR_PRINT`
  Neither the PRINT statement nor a WRITE statement with an asterisk (*) in place of a unit number includes an explicit logical unit number. Instead, both use an implicit internal logical unit number and the `FOR_PRINT` environment variable. If `FOR_PRINT` is not defined, the code `PRINT f,iolist` or `WRITE (*,f) iolist` writes to `stdout` (standard output). If `FOR_PRINT` is defined (as a filename optionally containing a path), the specified file would be written to.
- `FOR_READ`
  A READ statement that uses an asterisk (*) in place of a unit number does not include an explicit logical unit number. Instead, it uses an implicit internal logical unit number and the `FOR_READ` environment variable. If `FOR_READ` is not defined, the code `READ (*,f) iolist` or `READ f,iolist` reads from `stdin` (standard input). If `FOR_READ` is defined (as a filename optionally containing a path), the specified file would be read.
- `FOR_TYPE`
  The TYPE statement does not include an explicit logical unit number. Instead, it uses an implicit internal logical unit number and the `FOR_TYPE` environment variable. If `FOR_TYPE` is not defined, the code `TYPE f,iolist` writes to `stdout` (standard output). If `FOR_TYPE` is defined (as a filename optionally containing a path), the specified file would be written to.
- `FORT_BUFFERED`
  Lets you request that buffered I/O should be used at run time for output of all Fortran I/O units, except those with output to the terminal. This provides a run-time mechanism to support the `-assume buffered_io` `compiler option`.
- `FORT_CONVERTn`
  Lets you specify the data format for an unformatted file associated with a

particular unit number (*n*). See Methods of Specifying the Data Format: Overview and Environment Variable FORT_CONVERTn Method.

- `FORT_CONVERT.`*ext* and `FORT_CONVERT_`*ext*
  Lets you specify the data format for unformatted files with a particular file extension suffix (`ext`). See  Methods of Specifying the Data Format: Overview and Environment Variable FORT_CONVERT.ext or FORT_CONVERT_ext Method.
- `FORT`*n*
  Lets you specify the file name for a particular unit number (*n*), when a file name is not specified in the OPEN statement or an implicit OPEN is used, and the compiler option `-fpscomp filesfromcmd` was not specified. Preconnected files attached to units 0, 5, and 6 are by default associated with system standard I/O files.
- `NLSPATH`
  The path for the Intel Fortran run-time error message catalog.
- `TBK_ENABLE_VERBOSE_STACK_TRACE`
  This variable displays more detailed call stack information in the event of an error.
  The default brief output is usually sufficient to determine where an error occurred. Brief output includes up to twenty stack frames, reported one line per stack frame. For each frame, the image name containing the PC, routine name, line number, and source file are given.
  The verbose output, if selected, will provide (in addition to the information in brief output) the exception context record if the error was a machine exception (machine register dump), and for each frame, the return address, frame pointer and stack pointer and possible parameters to the routine. This output can be quite long (but limited to 16K bytes) and use of the environment variable `FOR_DIAGNOSTIC_LOG_FILE` is recommended if you want to capture the output accurately. Most situations should not require the use of verbose output.
  The variable `FOR_ENABLE_VERBOSE_STACK_TRACE` is also recognized for compatibility with Compaq* Fortran.
- `TBK_FULL_SRC_FILE_SPEC`
  This variable displays complete file name information for traceback output, including the path.
  By default, the traceback output displays only the file name and extension in the source file field. You must set this variable to display more.
  The variable `FOR_FULL_SRC_FILE_SPEC` is also recognized for compatibility with Compaq* Fortran.
- `TMP`, `TMPDIR`, and `TEMP`
  Specifies an alternate working directory where temporary files are created. See Temporary Files Created by the Compiler or Linker.

# Key IA-32 Compiler Files Summary

The following table shows files that are installed for use by the IA-32 compiler in `/opt/intel_fc_80/bin`:

| File | Description |
|---|---|
| `codecov` | Executable for the Code-coverage tool |
| `fortcom` | Executable used by the compiler |
| `fpp` | Fortran Preprocessor |
| `ifc` | For compatibility with previous releases |
| `ifc.cfg` | For compatibility with previous releases |
| `ifort` | Intel Fortran Compiler Version 8 |
| `ifortbin` | Executable used by the compiler |
| `ifort.cfg` | Configuration file |
| `ifortvars.csh` | Setup file for C shell |
| `ifortvars.sh` | Setup file for `bash` shell |
| `profmerge` | Utility used for Profile Guided Optimizations |
| `proforder` | Utility used for Profile Guided Optimizations |
| `tselect` | Test-prioritization tool |
| `uninstall.sh` | Uninstall utility |
| `xiar` | Tool used for Interprocedural Optimizations |
| `xild` | Tool used for Interprocedural Optimizations |

For a list of the files installed in `/lib`, see Libraries Provided by Intel Fortran.

# Key Itanium®-Based Compiler Files Summary

The following table shows files that are installed for use by the Itanium® -based compiler in `/opt/intel_fc_80/bin`:

| File | Description |
|---|---|
| `codecov` | Executable for the Code-coverage tool |
| `efc` | For compatibility with previous releases |
| `efc.cfg` | For compatibility with previous releases |
| `efcbin` | For compatibility with previous releases |
| `fortcom` | Executable used by the compiler |
| `fpp` | Fortran Preprocessor |
| `ias` | Intel assembler |
| `ifort` | Intel Fortran Compiler Version 8 |
| `ifort.cfg` | Configuration file |

| | |
|---|---|
| `ifortbin` | Executable used by the compiler |
| `ifortvars.csh` | Setup file for C shell |
| `ifortvars.sh` | Setup file for `bash` shell |
| `profdcg` | |
| `profmerge` | Utility used for Profile Guided Optimizations |
| `proforder` | Utility used for Profile Guided Optimizations |
| `tselect` | Test-prioritization tool |
| `uninstall.sh` | Uninstall utility |
| `xiar` | Tool used for Interprocedural Optimizations |
| `xild` | Tool used for Interprocedural Optimizations |

For a list of files installed in `/lib`, see Libraries Provided by Intel Fortran.

# Compiler Limits

The amount of data storage, the size of arrays, and the total size of executable programs are limited only by the amount of process virtual address space available, as determined by system parameters.

The table below shows the limits to the size and complexity of a single Intel Fortran program unit and to individual statements contained within it:

| Language Element | Limit |
|---|---|
| Actual number of arguments per CALL or function reference | Limited only by memory constraints |
| Arguments in a function reference in a specification expression | 255 |
| Array dimensions | 7 |
| Array construction nesting | 20 |
| Array elements per dimension | 9,223,372,036,854,775,807 = 2**31-1 on IA-32 systems; 2**63-1 Itanium-based systems; plus limited by current memory configuration |
| Constants: character and Hollerith | 7198 characters |
| Constants: characters read in list-directed I/O | 2048 characters |
| Continuation lines | 511 |
| Data and I/O implied DO nesting | 7 |
| DO and block IF statement nesting (combined) | 128 |
| DO loop index variable | 9,223,372,036,854,775,807 = 2**63- |

| | 1 |
|---|---|
| Format group nesting | 8 |
| Format statement length | 2048 characters |
| Fortran source line length | fixed form: 72 (or 132 if `-extend_source` is in effect) characters; free form: 7200 characters |
| INCLUDE file nesting | 20 levels |
| Labels in computed or assigned GOTO list | Limited only by memory constraints |
| Lexical tokens per statement | 20000 |
| Named common blocks | Limited only by memory constraints |
| Nesting of array constructor implied DOs | 7 |
| Nesting of input/output implied DOs | 7 |
| Nesting of interface blocks | Limited only by memory constraints |
| Nesting of DO, IF, or CASE constructs | Limited only by memory constraints |
| Nesting of parenthesized formats | Limited only by memory constraints |
| Number of digits in a numeric constant | Limited only by memory constraints |
| Parentheses nesting in expressions | Limited only by memory constraints |
| Structure nesting | 30 |
| Symbolic name length | 63 characters |

See the product Release Notes for more information on memory limits for large data objects.

# Hexadecimal-Binary-Octal-Decimal Conversions

The following table lists hexadecimal, binary, octal, and decimal conversion:

| Hex Number | Binary Number | Octal Number | Decimal Number |
|---|---|---|---|
| 0 | 0000 | 00 | 0 |
| 1 | 0001 | 01 | 1 |
| 2 | 0010 | 02 | 2 |
| 3 | 0011 | 03 | 3 |
| 4 | 0100 | 04 | 4 |
| 5 | 0101 | 05 | 5 |

| 6 | 0110 | 06 | 6 |
| 7 | 0111 | 07 | 7 |
| 8 | 1000 | 10 | 8 |
| 9 | 1001 | 11 | 9 |
| A | 1010 | 12 | 10 |
| B | 1011 | 13 | 11 |
| C | 1100 | 14 | 12 |
| D | 1101 | 15 | 13 |
| E | 1110 | 16 | 14 |
| F | 1111 | 17 | 15 |

# Compatibility with Previous Versions of Intel® Fortran

This topic is written for developers who are familiar with Intel Fortran Version 7.1 or earlier versions are now using Intel Fortran Version 8.

Intel® Fortran supports extensions to the ISO and ANSI standards, including a number of extensions defined by:

- Intel Fortran for various platforms
- Microsoft* Fortran PowerStation 4.0

Many language extensions associated with Microsoft* Fortran PowerStation Version 4 have been added to Intel Fortran.

## Differences Between Intel Fortran Version 7.1 and Intel Fortran Version 8

Some differences are:

- The command name for command-line use is now `ifort`. Earlier versions of Intel Fortran used a command name of `ifc` or `efc`. For Intel Fortran 8.0, these command names will still be accepted, but in some future Intel Fortran release, only the `ifort` command name will be accepted.
- The default configuration file name is now `ifort.cfg` instead of `ifl.cfg` or `efl.cfg`.

- The predefined symbol name for the Intel Fortran compiler is __INTEL_COMPILER and it has a value of 800.
- The record length (RECL specifier) for unformatted files is now 32-bit words. To get the record length in bytes, use the `/assume:byterecl` option.
- The backslash character (\) is not treated as an escape character for control sequences in character literals. To force the backslash to start escape sequences, use the `/assume:bscc` option.
- Intel Fortran Version 8 by default uses the integer -1 for the value of .TRUE. whereas Version 7 uses the integer 1 for the value of .TRUE.. If you use the `-fpscomp logicals` option with Version 8, the compiler will use the integer 1 for the value of .TRUE..
  Version 8 always uses the integer 0 as the value of .FALSE., as did Version 7.
  User-written routines in Fortran or other languages (for example, C) need to insure that they use values for .TRUE. and .FALSE. consistent with the compiler's choice.
- The random number generator used in Version 8 is different from the random number generator used in Version 7.
  Version 8 uses the random number generator based on the algorithm of Park and Miller, which is the generator used by Compaq Fortran. Version 7 used the Marsaglia random number generator.
  Both of these random number generators are compatible with the Fortran 90 standard.
  In addition, Version 8 uses different algorithms for the RANDOM_NUMBER and RANDOM_SEED intrinsics (compared to Version 7) and different algorithms are used for these intrinsics on IA32 and Itanium-based systems.

## Documentation Information

Some documentation has been moved. In particular:

- The *Intel Fortran User's Guide* now has separate parts for Building Applications and Optimizing Applications.
- Intel Fortran language information previously described in the *Intel Fortran Programmer's Reference*, including intrinsics procedures and directives, is now described in the online *Language Reference.*
- All Intel Fortran language elements and library routines are described in this online help file, allowing easy lookup of reference information.

## Version 7.1 Features Not Available in Intel Visual Fortran Version 8

The following Intel Fortran Version 7.1 features are not  available in Intel Visual Fortran Version 8:

- IMPLICIT AUTOMATIC | STATIC statements
- The Intel Fortran 8.0 run-time library system's ability to work with the Itanium processor simulator

# Run-Time Error Messages

The table below lists the errors processed by the Intel Fortran run-time library (RTL). For each error, the table provides the error number, the severity code, error message text, condition symbol name, and a detailed description of the errors.

To define the condition symbol values (PARAMETER statements) in your program, include the following file:

`/opt/intel_fc_80/include/for_iosdef.f`

As described in the table, the severity of the message determines which of the following occurs:: program execution continues with `info` and `warning`, the results might be incorrect with `error`, and program execution stops (unless a recovery method is specified) with `severe`. In the last case, to prevent program termination, you must include either an appropriate I/O error-handling specifier and recompile or, for certain errors, change the default action of a signal before you run the program again.

The first column lists error numbers returned to IOSTAT variables when an I/O error is detected.

The first line of the second column provides the message as it is displayed (following `forrtl:`), including the severity level, message number, and the message text. The following lines of the second column contain the status condition symbol (such as FOR$IOS_INCRECTYP) and an explanation of the message.

| Number | Severity Level, Number, and Message Text; Condition Symbol and Explanation |
|---|---|
| None [1] | info: Fortran error message number is *nnn*<br><br>The Intel Fortran message catalog file was not found on this system. See Message Catalog File Location. This error has no condition symbol. |

None [1] `warning: Could not open message catalog: for_msg.cat`

The Intel Fortran message catalog file was not found on this system. See Message Catalog File Location. This error has no condition symbol.

None [1] `info: Check environment variable NLSPATH and protection of` `pathname/for_msg.cat`

The Intel Fortran message catalog file was not found. See Message Catalog File Location. This error has no condition symbol.

None [1] `Insufficient memory to open Fortran RTL catalog: message 41`

The Intel Fortran message catalog file could not be opened because of insufficient virtual memory. To overcome this problem, increase the per-process data limit by using the `limit` (C shell) or `ulimit` (Bourne and Korn and bash  shells) commands before running the program again.

For more information, see error 41. This error has no condition symbol.

1 [1] `severe (1): Not a Fortran-specific error`

FOR$IOS_NOTFORSPE. An error in the user program or in the RTL was not a Intel Fortran-specific error and was not reportable through any other Intel Fortran run-time messages. If you call ERRSNS, an error of this kind returns a value of 1 (for more information on the ERRSNS subroutine, see the *Intel Fortran Language Reference Manual*).

8 `severe (8): Internal consistency check failure`

FOR$IOS_BUG_CHECK. Internal error. Please check that the program is correct. Recompile if an error existed in the program. If this error persists, submit a problem report.

9 `severe (9): Permission to access file denied`

FOR$IOS_PERACCFIL. Check the mode (protection) of the specified file. Make sure the correct file was being accessed. Change the protection, specified file, or process used before rerunning program.

10 `severe (10): Cannot overwrite existing file`

FOR$IOS_CAOVEEXI. Specified file *xxx* already exists when OPEN statement specified STATUS= ' NEW ' (create new file) using I/O unit *x*. Make sure correct file name, directory path, unit, and so forth were specified in the source program. Decide whether to:

- Rename or remove the existing file before rerunning the program.
- Modify the source file to specify different file specification, I/O unit, or OPEN statement STATUS.

**11** `info (11) `[1]`: Unit not connected`

FOR$IOS_UNINOTCON. The specified unit was not open at the time of the attempted I/O operation. Check if correct unit number was specified. If appropriate, use an OPEN statement to explicitly open the file (connect the file to the unit number).

**17** `severe (17): Syntax error in NAMELIST input`

FOR$IOS_SYNERRNAM. The syntax of input to a namelist-directed READ statement was incorrect.

**18** `severe (18): Too many values for NAMELIST variable`

FOR$IOS_TOOMANVAL. An attempt was made to assign too many values to a variable during a namelist READ statement.

**19** `severe (19): Invalid reference to variable in NAMELIST input`

FOR$IOS_INVREFVAR. One of the following conditions occurred:

- The variable was not a member of the namelist group.
- An attempt was made to subscript a scalar variable.
- A subscript of the array variable was out-of-bounds.
- An array variable was specified with too many or too few subscripts for the variable.
- An attempt was made to specify a substring of a noncharacter variable or array name.
- A substring specifier of the character variable was out-of-bounds.

- A subscript or substring specifier of the variable was not an integer constant.
- An attempt was made to specify a substring by using an unsubscripted array variable.

**20** `severe (20): REWIND error`

FOR$IOS_REWERR. One of the following conditions occurred:

- The file was not a sequential file.
- The file was not opened for sequential or append access.
- The Intel Fortran RTL I/O system detected an error condition during execution of a REWIND statement.

**21** `severe (21): Duplicate file specifications`

FOR$IOS_DUPFILSPE. Multiple attempts were made to specify file attributes without an intervening close operation. A DEFINE FILE statement was followed by another DEFINE FILE statement or an OPEN statement

**22** `severe (22): Input record too long`

FOR$IOS_INPRECTOO. A record was read that exceeded the explicit or default record length specified when the file was opened. To read the file, use an OPEN statement with a RECL= value (record length) of the appropriate size.

**23** `severe (23): BACKSPACE error`

FOR$IOS_BACERR. The Intel Fortran RTL I/O system detected an error condition during execution of a BACKSPACE statement.

**24** [1] `severe (24): End-of-file during read`

FOR$IOS_ENDDURREA. One of the following conditions occurred:

- A Intel Fortran RTL I/O system end-of-file condition was encountered during execution of a READ statement that did not contain an END, ERR, or IOSTAT specification.
- An end-of-file record written by the ENDFILE statement was encountered during execution of a

READ statement that did not contain an END,
ERR, or IOSTAT specification.

- An attempt was made to read past the end of an
internal file character string or array during
execution of a READ statement that did not
contain an END, ERR, or IOSTAT specification.

This error is returned by END and ERRSNS.

25 `severe (25): Record number outside range`

FOR$IOS_RECNUMOUT. A direct access READ,
WRITE, or FIND statement specified a record number
outside the range specified when the file was opened.

26 `severe (26): OPEN or DEFINE FILE required`

FOR$IOS_OPEDEFREQ. A direct access READ,
WRITE, or FIND statement was attempted for a file when
no prior DEFINE FILE or OPEN statement with
ACCESS= ' DIRECT ' was performed for that file.

27 `severe (27): Too many records in I/O statement`

FOR$IOS_TOOMANREC. An attempt was made to do
one of the following:

- Read or write more than one record with an
ENCODE or DECODE statement.
- Write more records than existed.

28 `severe (28): CLOSE error`

FOR$IOS_CLOERR. An error condition was detected by
the Intel Fortran RTL I/O system during execution of a
CLOSE statement.

29 `severe (29): File not found`

FOR$IOS_FILNOTFOU. A file with the specified name
could not be found during an open operation.

30 `severe (30): Open failure`

FOR$IOS_OPEFAI. An error was detected by the Intel
Fortran RTL I/O system while attempting to open a file in
an OPEN, INQUIRE, or other I/O statement. This
message is issued when the error condition is not one of
the more common conditions for which specific error
messages are provided. It can occur when an OPEN

operation was attempted for one of the following:

- Segmented file that was not on a disk or a raw magnetic tape
- Standard I/O file that had been closed

31 `severe (31): Mixed file access modes`

FOR$IOS_MIXFILACC. An attempt was made to use any of the following combinations:

- Formatted and unformatted operations on the same unit
- An invalid combination of access modes on a unit, such as direct and sequential
- A Intel Fortran RTL I/O statement on a logical unit that was opened by a program coded in another language

32 `severe (32): Invalid logical unit number`

FOR$IOS_INVLOGUNI. A logical unit number greater than 2,147,483,647 or less than zero was used in an I/O statement.

33 `severe (33): ENDFILE error`

FOR$IOS_ENDFILERR. One of the following conditions occurred:

- The file was not a sequential organization file with variable-length records.
- The file was not opened for sequential or append access.
- An unformatted file did not contain segmented records.
- The Intel Fortran RTL I/O system detected an error during execution of an ENDFILE statement.

34 `severe (34): Unit already open`

FOR$IOS_UNIALROPE. A DEFINE FILE statement specified a logical unit that was already opened.

35 `severe (35): Segmented record format error`

FOR$IOS_SEGRECFOR. An invalid segmented record control data word was detected in an unformatted

sequential file. The file was probably either created with RECORDTYPE=`'FIXED'` or `'VARIABLE'` in effect, or was created by a program written in a language other than Fortran.

36 `severe (36): Attempt to access non-existent record`

FOR$IOS_ATTACCNON. A direct-access READ or FIND statement attempted to access beyond the end of a relative file (or a sequential file on disk with fixed-length records) or access a record that was previously deleted in a relative file.

37 `severe (37): Inconsistent record length`

FOR$IOS_INCRECLEN. An attempt was made to open a direct access file without specifying a record length.

38 `severe (38): Error during write`

FOR$IOS_ERRDURWRI. The Intel Fortran RTL I/O system detected an error condition during execution of a WRITE statement.

39 `severe (39): Error during read`

FOR$IOS_ERRDURREA. The Intel Fortran RTL I/O system detected an error condition during execution of a READ statement.

40 `severe (40): Recursive I/O operation`

FOR$IOS_RECIO_OPE. While processing an I/O statement for a logical unit, another I/O operation on the same logical unit was attempted, such as a function subprogram that performs I/O to the same logical unit that was referenced in an expression in an I/O list or variable format expression.

41 `severe (41): Insufficient virtual memory`

FOR$IOS_INSVIRMEM. The Intel Fortran RTL attempted to exceed its available virtual memory while dynamically allocating space. To overcome this problem, increase the per-process data limit by using the `limit` (C shell) or `ulimit` (Bourne and Korn and bash shell) commands before you run this program again.

To determine whether the maximum per-process data size is already allocated, check the value of the *maxdsiz* parameter in the sysconfigtab or system configuration file. If necessary, increase its value. Changes to do not

take effect until the system has been rebooted (you do not need to rebuild the kernel if you modify sysconfigtab).

Before you try to run this program again, wait until the new system resources take effect.

**42** `severe (42): No such device`

FOR$IOS_NO_SUCDEV. A pathname included an invalid or unknown device name when an OPEN operation was attempted.

**43** `severe (43): File name specification error`

FOR$IOS_FILNAMSPE. A pathname or file name given to an OPEN or INQUIRE statement was not acceptable to the Intel Fortran RTL I/O system.

**44** `severe (44): Inconsistent record type`

FOR$IOS_INCRECTYP. The RECORDTYPE value in an OPEN statement did not match the record type attribute of the existing file that was opened.

**45** `severe (45): Keyword value error in OPEN statement`

FOR$IOS_KEYVALERR. An improper value was specified for an OPEN or CLOSE statement specifier requiring a value.

**46** `severe (46): Inconsistent OPEN/CLOSE parameters`

FOR$IOS_INCOPECLO. Specifications in an OPEN or CLOSE statement were inconsistent. Some invalid combinations follow:

- READONLY or ACTION= ' READ ' with STATUS= ' NEW ' or STATUS= ' SCRATCH '
- READONLY with STATUS= ' REPLACE ', ACTION= ' WRITE ', or ACTION= ' READWRITE '
- ACCESS= ' APPEND ' with READONLY, ACTION= ' READ ', STATUS= ' NEW ', or STATUS= ' SCRATCH '
- DISPOSE= ' SAVE ', ' PRINT ', or ' SUBMIT ' with STATUS= ' SCRATCH '
- DISPOSE= ' DELETE ' with READONLY
- CLOSE statement STATUS= ' DELETE ' with OPEN statement READONLY
- ACCESS= ' APPEND ' with STATUS= ' REPLACE '

248

- ACCESS='DIRECT' or 'KEYED' with
POSITION='APPEND', 'ASIS', or '
REWIND'

47 severe (47): Write to READONLY file

FOR$IOS_WRIREAFIL. A write operation was attempted
to a file that was declared ACTION='READ' or
READONLY in the OPEN statement that is currently in
effect.

48 severe (48): Invalid argument to Fortran Run-Time
Library

FOR$IOS_INVARGFOR. The compiler passed an invalid
or improperly coded argument to the Intel Fortran RTL.
This can occur if the compiler is newer than the RTL in
use.

51 severe (51): Inconsistent file organization

FOR$IOS_INCFILORG. The file organization specified in
an OPEN statement did not match the organization of the
existing file.

53 severe (53): No current record

FOR$IOS_NO_CURREC. Attempted to execute a
REWRITE statement to rewrite a record when the current
record was undefined. To define the current record,
execute a successful READ statement. You can
optionally perform an INQUIRE statement on the logical
unit after the READ statement and before the REWRITE
statement. No other operations on the logical unit may be
performed between the READ and REWRITE
statements.

55 severe (55): DELETE error

FOR$IOS_DELERR. An error condition was detected by
the Intel Fortran RTL I/O system during execution of a
DELETE statement.

57 severe (57): FIND error

FOR$IOS_FINERR. The Intel Fortran RTL I/O system
detected an error condition during execution of a FIND
statement.

58 [1] info (58): Format syntax error at or near *XX*

FOR$IOS_FMTSYN. Check the statement containing *xx*,
a character substring from the format string, for a format

syntax error. For information about FORMAT statements, see the *Intel Fortran Language Reference Manual*.

59 `severe (59): List-directed I/O syntax error`

FOR$IOS_LISIO_SYN [2]. The data in a list-directed input record had an invalid format, or the type of the constant was incompatible with the corresponding variable. The value of the variable was unchanged.

60 `severe (60): Infinite format loop`

FOR$IOS_INFFORLOO. The format associated with an I/O statement that included an I/O list had no field descriptors to use in transferring those values.

61 `severe or info` [3] `(61): Format/variable-type mismatch`

FOR$IOS_FORVARMIS [2]. An attempt was made either to read or write a real variable with an integer field descriptor (I, L, O, Z, B), or to read or write an integer or logical variable with a real field descriptor (D, E, or F).

62 `severe (62): Syntax error in format`

FOR$IOS_SYNERRFOR. A syntax error was encountered while the RTL was processing a format stored in an array or character variable.

63 `error or info` [3] `(63): Output conversion error`

FOR$IOS_OUTCONERR [2]. During a formatted output operation, the value of a particular number could not be output in the specified field length without loss of significant digits. When this situation is encountered, the overflowed field is filled with asterisks to indicate the error in the output record. If no ERR address has been defined for this error, the program continues after the error message is displayed.

64 `severe (64): Input conversion error`

FOR$IOS_INPCONERR [2]. During a formatted input operation, an invalid character was detected in an input field, or the input value overflowed the range representable in the input variable. The value of the variable was set to zero.

65 `error (65): Floating invalid`

FOR$IOS_FLTINV. During an arithmetic operation, the floating-point values used in a calculation were invalid for the type of operation requested or invalid exceptional

values. For example, when requesting a log of the floating-point values 0.0 or a negative number. For certain arithmetic expressions, specifying the `-check nopower` option can suppress this message.

66 `severe (66): Output statement overflows record`

FOR$IOS_OUTSTAOVE. An output statement attempted to transfer more data than would fit in the maximum record size.

67 `severe (67): Input statement requires too much data`

FOR$IOS_INPSTAREQ. Attempted to read more data than exists in a record with an unformatted READ statement or with a formatted sequential READ statement from a file opened with a PAD specifier value of `'NO'`.

68 `severe (68): Variable format expression value error`

FOR$IOS_VFEVALERR [2]. The value of a variable format expression was not within the range acceptable for its intended use; for example, a field width was less than or equal to zero. A value of 1 was assumed, except for a P edit descriptor, for which a value of zero was assumed.

69 [1] `error (69): Process interrupted (SIGINT)`

FOR$IOS_SIGINT. The process received the signal SIGINT. Determine source of this interrupt signal (described in `signal(3)`).

70 [1] `severe (70): Integer overflow`

FOR$IOS_INTOVF. During an arithmetic operation, an integer value exceeded byte, word, or longword range. The result of the operation was the correct low-order part. Consider specifying a larger integer data size (modify source program or, for an INTEGER declaration, possibly use the `f90` option `-integer_size` *nn* ).

71 [1] `severe (71): Integer divide by zero`

FOR$IOS_INTDIV. During an integer arithmetic operation, an attempt was made to divide by zero. The result of the operation was set to the dividend, which is equivalent to division by 1.

72 [1] `error (72): Floating overflow`

FOR$IOS_FLTOVF. During an arithmetic operation, a

floating-point value exceeded the largest representable value for that data type.

73 ¹error (73): Floating divide by zero

FOR$IOS_FLTDIV. During a floating-point arithmetic operation, an attempt was made to divide by zero.

74 ¹error (74): Floating underflow

FOR$IOS_FLTUND. During an arithmetic operation, a floating-point value became less than the smallest finite value for that data type. Depending on the values of the -fpe *n* option, the underflowed result was either set to zero or allowed to gradually underflow.

75 ¹error (75): Floating point exception

FOR$IOS_SIGFPE. A floating-point exception occurred. Core dump file created. Possible causes include:

- Division by zero
- Overflow
- Invalid operation, such as subtraction of infinite values, multiplication of zero by infinity (without signs), division of zero by zero or infinity by infinity
- Conversion of floating-point to fixed-point format when an overflow prevents conversion

76 ¹error (76): IOT trap signal

FOR$IOS_SIGIOT. Core dump file created. Examine core dump for possible cause of this IOT signal.

77 ¹severe (77): Subscript out of range

FOR$IOS_SUBRNG. An array reference was detected outside the declared array bounds.

78 ¹error (78): Process killed (SIGTERM)

FOR$IOS_SIGTERM. The process received the signal SIGTERM. Determine source of this software termination signal (described in signal(3) ).

79 ¹error (79): Process quit (SIGQUIT)

FOR$IOS_SIGQUIT. The process received the signal SIGQUIT. Core dump file created. Determine source of this quit signal (described in signal(3) ).

95 ¹info (95): Floating-point conversion failed

FOR$IOS_FLOCONFAI. The attempted unformatted

read or write of nonnative floating-point data failed because the floating-point value:

- Exceeded the allowable maximum value for the equivalent native format and was set equal to infinity (plus or minus)
- Was infinity (plus or minus) and was set to infinity (plus or minus)
- Was invalid and was set to not a number (NaN)

Very small numbers are set to zero (0). This error could be caused by the specified nonnative floating-point format not matching the floating-point format found in the specified file.

Check the following:

- Whether the correct file was specified.
- Whether the record layout matches the format Intel Fortran is expecting.
- The ranges for the data being used.
- Whether the correct nonnative floating-point data format was specified.

96 info (96): F_UFMTENDIAN environment variable was ignored:erroneous syntax

FOR$IOS_UFMTENDIAN. Syntax for specifying whether little endian or big endian conversion is performed for a given Fortran unit was incorrect. Even though the program will run, the results might not be correct if you do not change the value of F_UFMTENDIAN. For correct syntax, see Environment Variable F_UFMTENDIAN Method.

108 severe (108): Cannot stat file

FOR$IOS_CANSTAFIL. Attempted stat operation on the indicated file failed. Make sure correct file and unit were specified.

120 severe (120): Operation requires seek ability

FOR$IOS_OPEREQSEE. Attempted an operation on a file that requires the ability to perform seek operations on that file. Make sure the correct unit, directory path, and file were specified.

138 [1]`severe (138): Array index out of bounds (SIGILL)`

FOR$IOS_BRK_RANGE. Break exception generated a SIGTRAP signal (described in `signal(3)`). Core dump file created.

The cause is an array subscript that is outside the dimensioned boundaries of that array.

Either recompile with the `-check bounds` option (perhaps with the `decfort_dump_flag` environment variable set) or examine the core dump file to determine the source code in error.

139 [1]`severe (139): Array index out of bounds for index` *nn* `(SIGILL)`

FOR$IOS_BRK_RANGE2. Break exception generated a SIGTRAP signal (described in `signal(3)` ). Core dump file created.

The cause is an array subscript that is outside the dimensioned boundaries of the array index *n*.

Either recompile with the `-check bounds` option (perhaps with the `decfort_dump_flag` environment variable set) or examine the core dump file to determine the source code in error.

140 [1]`severe (140): Floating inexact`

FOR$IOS_FLTINE. A floating-point arithmetic or conversion operation gave a result that differs from the mathematically exact result. This trap is reported if the rounded result of an IEEE operation is not exact.

144 [1]`severe (144): reserved operand`

FOR$IOS_ROPRAND. The Intel Fortran RTL encountered a reserved operand. Please report the problem to Intel..

145 [1]`severe (145): Assertion error`

FOR$IOS_ASSERTERR. The Intel Fortran RTL encountered an assertion error. Please report the problem to Intel..

146 [1]`severe (146): Null pointer error`

FOR$IOS_NULPTRERR. Attempted to use a pointer that

does not contain an address. Modify the source program, recompile, and relink.

147 [1]`severe (147): stack overflow`

FOR$IOS_STKOVF. The Intel Fortran RTL encountered a stack overflow while executing your program.

148 [1]`severe (148): String length error`

FOR$IOS_STRLENERR. During a string operation, an integer value appears in a context where the value of the integer is outside the permissible string length range.

Either recompile with the `-check bounds` option (perhaps with the `decfort_dump_flag` environment variable set) or examine the `core` file to determine the source code causing the error.

149 [1]`severe (149): Substring error`

FOR$IOS_SUBSTRERR. An array subscript is outside the dimensioned boundaries of an array.

Either recompile with the `-check bounds` option (perhaps with the `decfort_dump_flag` environment variable set) or examine the `core` file to determine the source code causing the error.

150 [1]`severe (150): Range error`

FOR$IOS_RANGEERR. An integer value appears in a context where the value of the integer is outside the permissible range.

151 [1]`severe (151): Allocatable array is already allocated`

FOR$IOS_INVREALLOC. An allocatable array must not already be allocated when you attempt to allocate it. You must deallocate the array before it can again be allocated.

152 [1]`severe (152): Unresolved contention for Intel Fortran RTL global resource`

FOR$IOS_RESACQFAI. Failed to acquire a Intel Fortran RTL global resource for a reentrant routine.

For a multithreaded program, the requested global resource is held by a different thread in your program.

For a program using asynchronous handlers, the

requested global resource is held by the calling part of the program (such as main program) and your asynchronous handler attempted to acquire the same global resource.

153 [1]`severe (153): Allocatable array or pointer is not allocated`

FOR$IOS_INVDEALLOC. A Fortran-90 allocatable array or pointer must already be allocated when you attempt to deallocate it. You must allocate the array or pointer before it can again be deallocated.

173 [1]`severe (173): A pointer passed to DEALLOCATE points to an array that cannot be deallocated`

FOR$IOS_INVDEALLOC2. A pointer that was passed to DEALLOCATE pointed to an explicit array, an array slice, or some other type of memory that could not be deallocated in a DEALLOCATE statement. Only whole arrays previous allocated with an ALLOCATE statement can be validly passed to DEALLOCATE.

174 [1]`severe (174): SIGSEGV,` *message-text*

FOR$IOS_SIGSEGV. One of two possible messages occurs for this error number:

- `severe (174): SIGSEGV, segmentation fault occurred`

  This message indicates that the program attempted an invalid memory reference. Check the program for possible errors.

- `severe (174): SIGSEGV, possible program stack overflow occurred`

  The following explanatory text also appears:
  `Program requirements exceed current stacksize resource limit.`

175 [1]`severe (175): DATE argument to DATE_AND_TIME is too short (LEN=n), required LEN=8`

FOR$IOS_SHORTDATEARG. The number of characters associated with the DATE argument to the DATE_AND_TIME intrinsic was shorter than the required length. You must increase the number of characters passed in for this argument to be at least 8 characters in length. Verify that the TIME and ZONE arguments also meet their minimum lengths.

176 [1] `severe (176): TIME argument to DATE_AND_TIME is too short (LEN=n), required LEN=10`

FOR$IOS_SHORTTIMEARG. The number of characters associated with the TIME argument to the DATE_AND_TIME intrinsic was shorter than the required length. You must increase the number of characters passed in for this argument to be at least 10 characters in length. Verify that the DATE and ZONE arguments also meet their minimum lengths.

177 [1] `severe(177): ZONE argument to DATE_AND_TIME is too short (LEN=n), required LEN=5`

FOR$IOS_SHORTZONEARG. The number of characters associated with the ZONE argument to the DATE_AND_TIME intrinsic was shorter than the required length. You must increase the number of characters passed in for this argument to be at least 5 characters in length. Verify that the DATE and TIME arguments also meet their minimum lengths.

178 [1] `severe(178): Divide by zero`

FOR$IOS_DIV. A floating-point or integer divide-by-zero exception occurred.

179 [1,4] `severe(179): Cannot allocate array---overflow on array size calculation`

FOR$IOS_ARRSIZEOVF. An attempt to dynamically allocate storage for an array failed because the required storage size exceeds addressable memory.

256 `severe (256): Unformatted I/O to unit open for formatted transfers`

FOR$IOS_UNFIO_FMT. Attempted unformatted I/O to a unit where the OPEN statement (FORM specifier) indicated the file was formatted. Check that the correct unit (file) was specified.

If the FORM specifier was not present in the OPEN statement and the file contains unformatted data, specify FORM='UNFORMATTED' in the OPEN statement. Otherwise, if appropriate, use formatted I/O (such as list-directed or namelist I/O).

257 `severe (257): Formatted I/O to unit open for unformatted transfers`

FOR$IOS_FMTIO_UNF. Attempted formatted I/O (such as list-directed or namelist I/O) to a unit where the OPEN

statement indicated the file was unformatted (FORM specifier). Check that the correct unit (file) was specified.

If the FORM specifier was not present in the OPEN statement and the file contains formatted data, specify FORM= ' FORMATTED ' in the OPEN statement. Otherwise, if appropriate, use unformatted I/O.

264 `severe (264): operation requires file to be on disk or tape`

FOR$IOS_OPERREQDIS. Attempted to use a BACKSPACE statement on such devices as a terminal or pipe.

265 `severe (265): operation requires sequential file organization and access`

FOR$IOS_OPEREQSEQ. Attempted to use a BACKSPACE statement on a file whose organization was not sequential or whose access was not sequential. A BACKSPACE statement can only be used for sequential files opened for sequential access.

266 [1]`error (266): Fortran abort routine called`

FOR$IOS_PROABOUSE. The program called `abort` to terminate the program.

268 [1]`severe (268): End of record during read`

FOR$IOS_ENDRECDUR. An end-of-record condition was encountered during execution of a nonadvancing I/O READ statement that did not specify the EOR branch specifier.

297 [1]`info (297): nn floating invalid traps`

FOR$IOS_FLOINVEXC. The total number of floating-point invalid data traps encountered during program execution was $nn$. This message appears at program completion.

298 [1]`info (298): nnfloating overflow traps`

FOR$IOS_FLOOVFEXC. The total number of floating-point overflow traps encountered during program execution was $nn$. This message appears at program completion.

299 [1]`info (299): nnfloating divide-by-zero traps`

FOR$IOS_FLODIV0EXC. The total number of floating-point divide-by-zero traps encountered during program

execution was *nn*. This message appears at program completion.

300 [1] `info (300): `*nn*`floating underflow traps`

FOR$IOS_FLOUNDEXC. The total number of floating-point underflow traps encountered during program execution was *nn*. This message appears at program completion.

**Footnotes:**

1 Identifies errors not returned by IOSTAT.

2 The ERR transfer is taken after completion of the I/O statement for error numbers 59, 61, 63, 64, and 68. The resulting file status and record position are the same as if no error had occurred. However, other I/O errors take the ERR transfer as soon as the error is detected, so file status and record position are undefined.

3 For errors 61 and 63, the severity depends on the `-check` options used during compilation.

4 Identifies errors that can be returned by STAT in an ALLOCATE statement.

# Index