Intel® Fortran Composer XE 2011 for Linux*
Installation Guide and Release Notes

Document number: 321415-003US
3 February 2012

Table of Contents

1 Introduction 3
1.1 Change History 3
1.2 Product Contents 5
1.3 System Requirements 5
1.3.1 Red Hat Enterprise Linux* 4 Support Deprecated 6
1.3.2 Asianux* Support Deprecated 6
1.3.3 IA-64 Architecture (Intel® Itanium®) Development Not Supported 6
1.4 Documentation 6
1.5 Japanese Language Support 7
1.6 Technical Support 7
2 Installation 7
2.1 Cluster Installation 8
2.2 Activation of Purchase after Evaluation Using the Intel Activation Tool 8
2.3 Silent Install 8
2.4 Using a License Server 8
2.5 Known Installation Issues 8
2.6 Installation Folders 9
2.7 Removal/Uninstall 10
3 Intel® Fortran Compiler 10
3.1 Compatibility 10
3.1.1 Stack Alignment Change for REAL(16) and COMPLEX(16) Datatypes 10
3.2 New and Changed Features 11
3.2.1 Features from Fortran 2003 11
3.2.2 Features from Fortran 2008 11
3.2.3 Coarrays 12

3.2.4 Static Security Analysis Feature (formerly Source Checker) Requires Intel® Inspector XE. 14
3.2.5 New and Changed Directives (Update 6) 14
3.2.6 OpenMP Changes (Update 6) 14
3.2.7 Other Changes 14
3.3 New and Changed Compiler Options 15
3.3.1 New and Changed in Composer XE 2011 Update 6 (and later) 15
3.3.2 New and Changed in Composer XE 2011 (through Update 5) 16
3.3.3 –assume [no]std_intent_in added (Update 9) 16
3.3.4 Additional Keywords for `-sox` option, default changed (Update 3) 17
3.4 Other Changes and Notes 17
3.4.1 Optimization Reports Disabled by Default 17
3.4.2 Establishing the Compiler Environment 17
3.4.3 OpenMP* Legacy Libraries Removed 17
3.4.4 RANF Portability Function Is Now an Intrinsic 17
Known Issues 17
3.5 17
3.5.1 Coarray Issues 17
3.5.2 Allocatable Arrays and OpenMP* PRIVATE or FIRSTPRIVATE 18
3.6 Fortran 2003 and Fortran 2008 Feature Summary 18
4 Intel® Debugger (IDB) 21
4.1 Setting up the Java* Runtime Environment 21
4.2 Starting the Debugger 21
4.3 Additional Documentation 21
4.4 Debugger Features 21
4.4.1 Changes for Intel® Fortran Composer XE 2011 Update 6 21
4.4.2 Main Features of IDB 22
4.5 Known Issues and Changes 23
4.5.1 Coarray elements cannot be viewed. 23
4.5.2 Signals Dialog not working Signals Dialog not working 23
4.5.3 Resizing GUI 23
4.5.4 `$cdir, $cwd` Directories 23
4.5.5 `info stack` Usage 23
4.5.6 `$stepg0` Default Value Changed 24
4.5.7 SIGTRAP error on some Linux* Systems 24
4.5.8 idb GUI cannot be used to debug MPI processes 24

4.5.9 Thread Syncpoint Creation in GUI 24
4.5.10 Stack Alignment for IA-32 Architecture 24
4.5.11 GNOME Environment Issues 24
4.5.12 Accessing Online-Help 24
4.5.13 Debugger crashes if $HOME not set on calling shell 25
4.5.14 Command line parameter –parallel not supported 25
4.5.15 Command line parameter –idb and -dbx not supported 25
4.5.16 Core File Debugging 25
4.5.17 Thread Data Sharing Filters may not work correctly 25
4.5.18 Default .gdbinit script on Pardus systems may cause the Debugger crash 25
4.5.19 No thread info available on Pardus systems 25
5 Intel® Math Kernel Library 25
5.1 What's New in Intel® MKL 10.3 Update 9 25
5.2 What's New in Intel® MKL 10.3 Update 8 26
5.3 What's New in Intel® MKL 10.3 Update 7 26
5.4 What's New in Intel® MKL 10.3 Update 6 26
5.5 What's New in Intel® MKL 10.3 Update 5 27
5.6 What's New in Intel® MKL 10.3 Update 4 27
5.7 What's New in Intel® MKL 10.3 Update 3 27
5.8 What's New in Intel® MKL 10.3 Update 2 28
5.9 What's New in Intel® MKL 10.3 Update 1 28
5.10 What's New in Intel® MKL 10.3 29
5.11 Known Issues 30
5.12 Notices 30
5.13 Attributions 31
6 Disclaimer and Legal Information 31

1 Introduction

This document describes how to install the product, provide a summary of new and changed product features and includes notes about features and problems not described in the product documentation. Intel® Fortran Composer XE 2011 is the next release of the product formerly called Intel® Fortran Compiler Professional Edition.

1.1 Change History

This section highlights important changes in product updates to Intel® Fortran Composer XE 2011 for Linux.

Update 9

- Intel® Fortran Compiler updated to 12.1.3
  - –assume std_intent_in option added
  - Restriction on ALLOCATABLE or POINTER components of derived type coarray noted
- Intel® Math Kernel Library updated to 10.3 Update 9
- Corrections to reported problems

Update 8

- Intel® Fortran Compiler updated to 12.1.2
- Intel® Math Kernel Library updated to 10.3 Update 8
- Pardus 2011* added as a supported distribution
- Corrections to reported problems

Update 7

- Intel® Fortran Compiler updated to 12.1.1
- Intel® Math Kernel Library updated to 10.3 Update 7
  - Intel® MKL Notices updated to include more deprecations
- Corrections to reported problems

Update 6

- The product installs into a new top-level folder.
- Cluster installation is now supported
- Intel® Fortran Compiler updated to 12.1.0
  - Additional Fortran 2003 and Fortran 2008 features supported
  - Additional compiler options
  - Additions and enhancements to general directives
  - Enhancements to OpenMP* support
  - The core compiler documentation known as the User and Reference Guides has been reorganized and streamlined. Among the most noticeable changes are: a new Key Features section highlighting important Intel compiler functionality and the organization of the Compiler Option reference section into functional groups.
- Intel® Math Kernel Library updated to 10.3 Update 6
- Corrections to reported problems

Update 5

- Intel® Fortran Compiler updated to 12.0.5
- Intel® Math Kernel Library updated to 10.3 Update 5
- Support for Asianux* is deprecated
- Corrections to reported problems

Update 4

- Intel® Fortran Compiler updated to 12.0.4
- Intel® Math Kernel Library updated to 10.3 Update 4
- Corrections to reported problems

Update 3

- Intel® Fortran Compiler updated to 12.0.3
  - The `-sox` option now accepts optional keywords specifying information to include and the default behavior has changed
- Japanese localized documentation and diagnostic messages now available
- Intel® Math Kernel Library updated to 10.3 Update 3
- Corrections to reported problems

Update 2

- Intel® Fortran Compiler updated to 12.0.2
- Intel® Math Kernel Library updated to 10.3 Update 2
- The way that the Static Security Analysis feature creates data files has changed
- Corrections to reported problems

Update 1

- Intel® Fortran Compiler updated to 12.0.1
- Intel® Math Kernel Library updated to 10.3 Update 1
- Corrections to reported problems

Intel® Fortran Composer XE 2011 Product Release

- Initial product release

1.2 Product Contents

*Intel® Fortran Composer XE 2011 for Linux** includes the following components:

- Intel® Fortran Compiler XE 12.1.3 for building applications that run on IA-32 and Intel® 64 architecture systems running the Linux* operating system
- Intel® Debugger 12.1.3
- Intel® Math Kernel Library 10.3 Update 9
- On-disk documentation

1.3 System Requirements

For an explanation of architecture names, see http://intel.ly/mXIljK

- A PC based on an IA-32 or Intel® 64 architecture processor supporting the Intel® Streaming SIMD Extensions 2 (Intel® SSE2) instructions (Intel® Pentium® 4 processor or later, or compatible non-Intel processor)
    - Development of 64-bit applications is supported on a 64-bit version of the OS only. Development of 32-bit applications is supported on either 32-bit or 64-bit versions of the OS.
    - Development for a 32-bit on a 64-bit host may require optional library components (ia32-libs, lib32gcc1, lib32stdc++6, libc6-dev-i386, gcc-multilib) to be installed from your Linux distribution.
- For the best experience, a multi-core or multi-processor system is recommended
- 1GB of RAM (2GB recommended)
- 2GB free disk space for all features
- One of the following Linux distributions (this is the list of distributions tested by Intel; other distributions may or may not work and are not recommended - please refer to Technical Support if you have questions):
    - Asianux* 3.0, 4.0 (Deprecated)
    - Debian* 6.0
    - Fedora* 15
    - Red Hat Enterprise Linux* 4 (Deprecated), 5, 6
    - SUSE LINUX Enterprise Server* 10,11 SP1
    - Ubuntu* 10.04, 11.04
    - Intel® Cluster Ready
    - Pardus* 2011.2 (x64 only)
- Linux Developer tools component installed, including gcc, g++ and related tools

- Library `libunwind.so` is required in order to use the `-traceback` option. Some Linux distributions may require that it be obtained and installed separately.

Additional requirements to use the Graphical User Interface of the Intel® Debugger

- IA-32 Architecture system or Intel® 64 Architecture system
- Java* Runtime Environment (JRE) 5.0 (also called 1.5)
- A 32-bit JRE must be used on an IA-32 architecture system and a 64-bit JRE must be used on an Intel® 64 architecture system

Notes

- The Intel compilers are tested with a number of different Linux distributions, with different versions of gcc. Some Linux distributions may contain header files different from those we have tested, which may cause problems. The version of glibc you use must be consistent with the version of gcc in use. For best results, use only the gcc versions as supplied with distributions listed above.
- The default for the Intel® compilers is to build IA-32 architecture applications that require a processor supporting the Intel® SSE2 instructions - for example, the Intel® Pentium® 4 processor. A compiler option is available to generate code that will run on any IA-32 architecture processor.
- Compiling very large source files (several thousands of lines) using advanced optimizations such as -O3, -ipo and -openmp, may require substantially larger amounts of RAM.
- The above lists of processor model names are not exhaustive - other processor models correctly supporting the same instruction set as those listed are expected to work. Please refer to Technical Support if you have questions regarding a specific processor model
- Some optimization options have restrictions regarding the processor type on which the application is run. Please see the documentation of these options for more information.

1.3.1 Red Hat Enterprise Linux* 4 Support Deprecated

In a future major release of Intel® Fortran Composer XE, support will be removed for installation and use on Red Hat Enterprise Linux 4. Intel recommends migrating to a newer version of this operating systems.

1.3.2 Asianux* Support Deprecated

In a future major release of Intel® Fortran Composer XE, support will be removed for installation and use on all Asianux distributions.

1.3.3 IA-64 Architecture (Intel® Itanium®) Development Not Supported

This product version does not support development on or for IA-64 architecture (Intel® Itanium®) systems. The version 11.1 compiler remains available for development of IA-64 architecture applications.

1.4 Documentation

Product documentation can be found in the `Documentation` folder as shown under Installation Folders.

Optimization Notice
Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel

microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.
Notice revision #20110804

1.5 Japanese Language Support

Intel compilers optionally provide support for Japanese language users when the combined English-Japanese product is installed. Error messages, visual development environment dialogs and some documentation are provided in Japanese in addition to English. By default, the language of error messages and dialogs matches that of your operating system language selection. Japanese-language documentation can be found in the `ja_JP` subdirectory for documentation and samples.
Japanese language support is not provided with every update of the product.
If you wish to use Japanese-language support on an English-language operating system, or English-language support on a Japanese-language operating system, you will find instructions at http://intel.ly/pla2A5

1.6 Technical Support

Register your license at the Intel® Software Development Products Registration Center. Registration entitles you to free technical support, product updates and upgrades for the duration of the support term.
For information about how to find Technical Support, Product Updates, User Forums, FAQs, tips and tricks, and other support information, please visit: http://www.intel.com/software/products/support/
Note: If your distributor provides technical support for this product, please contact them for support rather than Intel.

2 Installation

The installation of the product requires a valid license file or serial number. If you are evaluating the product, you can also choose the "Evaluate this product (no serial number required)" option during installation
If you received your product on DVD, mount the DVD, change the directory (`cd`) to the top-level directory of the mounted DVD and begin the installation using the command:

```
./install.sh
```

If you received the product as a downloadable file, first unpack it into a writeable directory of your choice using the command:

```
tar -xzvf name-of-downloaded-file
```

Then change the directory (`cd`) to the directory containing the unpacked files and begin the installation using the command:

```
./install.sh
```

Follow the prompts to complete installation.
Note that there are several different downloadable files available, each providing different combinations of components. Please read the download web page carefully to determine which file is appropriate for you.
You do not need to uninstall previous versions or updates before installing a newer version – the new version will coexist with the older versions.

2.1 Cluster Installation

If a license for Intel® Cluster Studio XE is present, and the installation detects that the installing system is a member of a cluster, the product will be installed on all visible nodes of the cluster when a "Full" installation is requested. If a "Custom" installation is requested, you will be given the option to install on the current node only.

2.2 Activation of Purchase after Evaluation Using the Intel Activation Tool

Note for evaluation customers: the Intel Activation Tool "`Activate`" is included in this product release and installed at `/opt/intel/ActivationTool/Activation/` directory.
If you installed the product using an Evaluation license or serial number (SN), or using the "Evaluate this product (no serial number required)" option during installation, and then purchased the product, you can activate your purchase using the Intel Activation Tool at `/opt/intel/ActivationTool/Activation/Activate`. It will convert your evaluation software to a fully licensed product. To use the tool:

```
$ /opt/intel/ActivationTool/Activation/Activate [SN_here]
```

2.3 Silent Install

For information on automated or "silent" install capability, please see http://intel.ly/ngVHY8

2.4 Using a License Server

If you have purchased a "floating" license, see http://intel.ly/oPEdEe for information on how to install using a license file or license server. This article also provides a source for the Intel® License Manager for FLEXlm* product that can be installed on any of a wide variety of systems.

2.5 Known Installation Issues

- If you have enabled the Security-Enhanced Linux (SELinux) feature of your Linux distribution, you must change the `SELINUX` mode to `permissive` before installing the Intel Fortran Compiler. Please see the documentation for your Linux distribution for details. After installation is complete, you may reset the `SELINUX` mode to its previous value.
- On some versions of Linux, auto-mounted devices do not have the "exec" permission and therefore running the installation script directly from the DVD will result in an error such as:

  ```
  bash: ./install.sh: /bin/bash: bad interpreter: Permission denied
  ```

  If you see this error, remount the DVD with exec permission, for example:

  ```
  mount /media/<dvd_label> -o remount,exec
  ```

  and then try the installation again.
- The product is fully supported on Ubuntu and Debian Linux distributions for IA-32 and Intel® 64 architecture systems as noted above under System Requirements. Due to a restriction in the licensing software, however, it is not possible to use the Trial License feature when evaluating IA-32 components on an Intel® 64 architecture system under Ubuntu or Debian. This affects using a Trial License only. Use of serial numbers, license files, floating licenses or other license manager operations, and off-line activation (with serial numbers) is not affected. If you need to evaluate IA-32 components of the product on an Intel® 64 architecture Ubuntu or Debian system, please visit the Intel® Software Evaluation Center (http://intel.ly/nJS8y8) to obtain an evaluation serial number.

2.6 Installation Folders

The compiler installs, by default, under `/opt/intel` – this is referenced as `<install-dir>` in the remainder of this document. You are able to specify a different location, and can also perform a "non-root" install in the location of your choice.

The directory organization has changed since the Intel® Compilers 11.1 release.

While the top-level installation directory has also changed between the original Intel Fortran Composer XE 2011 release and Composer XE 2011 Update 6, the `composerxe` symbolic link can still be used to reference the latest product installation.

Under `<install-dir>` are the following directories:

- `bin` – contains symbolic links to executables for the latest installed version
- `lib` – symbolic link to the lib directory for the latest installed version
- `include` – symbolic link to the include directory for the latest installed version
- `man` – symbolic link to the directory containing man pages for the latest installed version
- `mkl` – symbolic link to the directory for the latest installed version of Intel® Math Kernel Library
- `composerxe` – symbolic link to the *`composer`*`_xe_2011_sp1` directory
- `composer_xe_2011_sp1` – directory containing symbolic links to subdirectories for the latest installed Intel® Composer XE 2011 product release
- `composer_xe_2011_sp1.<n>.<pkg>` - physical directory containing files for a specific update version. `<n>` is the update number, and `<pkg>` is a package build identifier

Each `composer_xe_2011_sp1` directory contains the following directories that reference the latest installed Intel® Composer XE 2011 product:

- `bin` – directory containing scripts to establish the compiler environment and symbolic links to compiler executables for the host platform
- `pkg_bin` – symbolic link to the compiler `bin` directory
- `include` – symbolic link to the compiler `include` directory
- `lib` – symbolic link to the compiler `lib` directory
- `mkl` – symbolic link to the `mkl` directory
- `debugger` – symbolic link to the `debugger` directory
- `man` – symbolic link to the directory containing man pages for the latest installed version
- `Documentation` – symbolic link to the `documentation` directory
- `Samples` – symbolic link to the `samples` directory
- `eclipse_support` – symbolic link to a directory created by the Intel Debugger component that is shared between Intel Fortran and Intel C++. Intel does not provide Eclipse support for Fortran.

Each `composer_xe_2011_sp1.<n>.<pkg>` directory contains the following directories that reference a specific update of the Intel® Composer XE 2011 compiler:

- `bin` – all executables
- `compiler` – shared libraries and include/header files
- `debugger` – debugger files
- `Documentation` – documentation files
- `eclipse_support` –directory created by the Intel Debugger component that is shared between Intel Fortran and Intel C++. Intel does not provide Eclipse support for Fortran.
- `man` – `man` pages
- `mkl` – Intel® Math Kernel Library libraries and header files
- `mpirt` – Intel® MPI Library run-time files used by Fortran coarray support
- `Samples` – Product samples and tutorial files

If you have both the Intel C++ and Intel Fortran compilers installed, they will share folders for a given version and update.
This directory layout allows you to choose whether you want the latest product update, no matter which version, the latest update of the Intel® Composer XE 2011 product, or a specific update. Most users will reference `<install-dir>/bin` for the `compilervars.sh [.csh]` script, which will always get the latest product installed. This layout should remain stable for future releases.

2.7 Removal/Uninstall

Removing (uninstalling) the product should be done by the same user who installed it (root or a non-root user). If `sudo` was used to install, it must be used to uninstall as well. It is not possible to remove the compiler while leaving any of the performance library components installed.

1. Open a terminal window and set default (`cd`) to any folder outside `<install-dir>`
2. Type the command: `<install-dir>/bin/uninstall.sh`
3. Follow the prompts
4. Repeat steps 2 and 3 to remove additional platforms or versions

If you also have the same-numbered version of Intel® C++ Compiler installed, it may also be removed.

3 Intel® Fortran Compiler

This section summarizes changes, new features and late-breaking news about the Intel Fortran Compiler.

3.1 Compatibility

In general, object code and modules compiled with earlier versions of Intel Fortran Compiler for Linux* (8.0 and later) may be used in a build with version 12. Exceptions include:

- Sources that use the CLASS keyword to declare polymorphic variables and which were built with a compiler version earlier than 12.0 must be recompiled.
- Objects built with the multi-file interprocedural optimization (-ipo) option must be recompiled.
- Objects that use the REAL(16) , REAL*16, COMPLEX(16) or COMPLEX*32 datatypes and which were compiled with versions earlier than 12.0 must be recompiled.
- Objects built for the Intel® 64 architecture with a compiler version earlier than 10.0 and that have module variables must be recompiled. If non-Fortran sources reference these variables, the external names may need to be changed to remove an incorrect leading underscore.
- Modules that specified an ATTRIBUTES ALIGN directive and were compiled with versions earlier than 11.0 must be recompiled. The compiler will notify you if this issue is encountered.

3.1.1 Stack Alignment Change for REAL(16) and COMPLEX(16) Datatypes

In releases prior to compiler version 12.0, when a REAL(16) or COMPLEX(16) (REAL*16 or COMPLEX*32) item was passed by value, the stack address was aligned at 4 bytes. For improved performance, the version 12.0 (and later) compiler aligns such items at 16 bytes and expects received arguments to be aligned on 16-byte boundaries. This change is also compatible with gcc.
This change primarily affects compiler-generated calls to library routines that do computations on REAL(16) values, including intrinsics. If you have code compiled with earlier versions and link it with the version 12.0 or libraries, or have an application linked to the shared version of the Intel run-time libraries, it may give incorrect results.
In order to avoid errors, you must recompile all Fortran sources that use the REAL(16) and COMPLEX(16) datatypes.

3.2 New and Changed Features

Some language features may not yet be described in the compiler documentation. Please refer to the Fortran 2003 Standard (http://j3-fortran.org/doc/2003_Committee_Draft/04-007.pdf) and Fortran 2008 Standard (http://j3-fortran.org/doc/standing/links/007.pdf) if necessary.

3.2.1 Features from Fortran 2003

- FINAL subroutines
- GENERIC keyword for type-bound procedures
- A generic interface may have the same name as a derived type
- Bounds specification and bounds remapping list on a pointer assignment
- ALLOCATE with SOURCE= (polymorphic source supported in Update 6)

3.2.2 Features from Fortran 2008

- Maximum array rank has been raised to 31 dimensions (Fortran 2008 specifies 15)
- Coarrays
- CODIMENSION attribute
- SYNC ALL statement
- SYNC IMAGES statement
- SYNC MEMORY statement
- CRITICAL and END CRITICAL statements
- LOCK and UNLOCK statements
- ERROR STOP statement
- ALLOCATE and DEALLOCATE may specify coarrays
- Intrinsic procedures IMAGE_INDEX, LCOBOUND, NUM_IMAGES, THIS_IMAGE, UCOBOUND
    - **Note:** ATOMIC_DEFINE and ATOMIC_REF are not supported in this version
- CONTIGUOUS attribute
- MOLD keyword in ALLOCATE
- DO CONCURRENT
- NEWUNIT keyword in OPEN
- G0 and G0.d format edit descriptor
- Unlimited format item repeat count specifier
- A CONTAINS section may be empty
- Intrinsic procedures BESSEL_J0, BESSEL_J1, BESSEL_JN, BESSEL_YN, BGE, BGT, BLE, BLT, DSHIFTL, DSHIFTR, ERF, ERFC, ERFC_SCALED, GAMMA, HYPOT, IALL, IANY, IPARITY, IS_CONTIGUOUS, LEADZ, LOG_GAMMA, MASKL, MASKR, MERGE_BITS, NORM2, PARITY, POPCNT, POPPAR, SHIFTA, SHIFTL, SHIFTR, STORAGE_SIZE, TRAILZ
- Additions to intrinsic module ISO_FORTRAN_ENV: ATOMIC_INT_KIND, ATOMIC_LOGICAL_KIND, CHARACTER_KINDS, INTEGER_KINDS, INT8, INT16, INT32, INT64, LOCK_TYPE, LOGICAL_KINDS, REAL_KINDS, REAL32, REAL64, REAL128, STAT_LOCKED, STAT_LOCKED_OTHER_IMAGE, STAT_UNLOCKED
- (Update 6) An OPTIONAL dummy argument that does not have the ALLOCATABLE or POINTER attribute, and which corresponds to an actual argument that: has the ALLOCATABLE attribute and is not allocated, or has the POINTER attribute and is disassociated, or is a reference to the intrinsic function NULL, is considered not present
- (Update 6) A dummy argument that is a procedure pointer may be associated with an actual argument that is a valid target for the dummy pointer, or is a reference to the intrinsic function NULL. If the actual argument is not a pointer, the dummy argument shall have the INTENT(IN)

attribute.

3.2.3 Coarrays

No special procedure is necessary to run a program that uses coarrays in a shared-memory configuration; you simply run the executable file. The underlying parallelization implementation is Intel® MPI. Installation of the compiler automatically installs the necessary Intel® MPI run-time libraries to run on shared memory. The Intel® Cluster Toolkit product (optional) installs the necessary Intel® MPI run-time libraries to run on distributed memory. Use of coarray applications with any other MPI implementation, or with OpenMP*, is not supported.

By default, the number of images created is equal to the number of execution units on the current system. You can override that by specifying the option `-coarray-num-images <n>` on the ifort command that compiles the main program. You can also specify the number of images in an environment variable `FOR_COARRAY_NUM_IMAGES`.

3.2.3.1 How to Debug a Coarray Application

The following instructions describe how to debug a Coarray application.

1. Add a stall loop to your application before the area of code you wish to debug, e.g.:

```
LOGICAL VOLATILE :: WAIT_FOR_DEBUGGER
LOGICAL, VOLATILE :: TICK
   :
DO WHILE(WAIT_FOR_DEBUGGER)
  TICK = .NOT. TICK
  END DO
! Code you want to debug is here
!
```

   The use of `VOLATILE` is required to ensure that the loop will not be removed by the compiler. If the problem is only found on one image, you can wrap the loop in `IF (THIS_IMAGE() .EQ. 4) THEN` or the like.
2. Compile and link with debug enabled (`-g`).
3. Create at least N+1 terminal windows on the machine where the application will be running, where N is the number of images your application will have.
4. In a terminal window, start the application.
   `linuxprompt> ./my_app`
5. In each of the other terminal windows, set your default directory to be the same as the location of the application executable. Use the ps command in one of the windows to find out which processes are running your application:

   ```
   linuxprompt> ps -ef | grep 'whoami' | grep my_app
   ```

   There will be several processes. The oldest is the one you started in step 4 – it has run the MPI launcher and is now waiting for the others to terminate. Do not debug it.

   The others will look like this:

   ```
   <your-user-name>  25653 25650 98 15:06 ?        00:00:49 my_app
   <your-user-name>  25654 25651 97 15:06 ?        00:00:48 my_app
   ```

```
<your-user-name>  25655 25649 98 15:06 ?        00:00:49 my_app
```

The first number is the PID of the process (e.g., 25653 in the first line).

Call the PIDs of these N processes running "my_app" P1, P2, P3 and so on.

6. In each window other than the first, start your debugger and set it to stop processes when attached:

```
linuxprompt> idb -idb
(idb) set $stoponattach = 1
```

or

```
linuxprompt> gdb
```

7. Attach to one of the processes (e.g. to P1 in window 1, to P2 in window 2, etc.)

```
(idb) attach <P1> my_app
```

or

```
(gdb) attach <P1>
```

8. Get execution out of the stall loop:

(idb) assign WAIT_FOR_DEBUGGER = .FALSE.

or

(gdb) set WAIT_FOR_DEBUGGER = .false.

9. You can now debug.

If you are using idb, you can use the multiprocess capability of idb to have only one debugger window instead of N. First, attach to each process and get out of the loop (steps 7 and 8).

```
(idb) attach <P1> my_app
(idb) assign WAIT_FOR_DEBUGGER = .FALSE.
(idb) attach <P2> my_app
(idb) assign WAIT_FOR_DEBUGGER = .FALSE.
(idb) attach <P3> my_app
(idb) assign WAIT_FOR_DEBUGGER = .FALSE.
```

Use the "process" command to switch debugging focus from one process to another:

```
(idb) process <Pn>
```

Processes not focused on will remain in the state they were left in: with breakpoints and watchpoints set but not running.

3.2.3.2 Coarray Known Issues

The following features are known not to work in this version:

- Output (WRITE, PRINT, etc.) of an array slice of a coarray referencing another image. A whole array reference, or a single element works.
- Default initialization of a REAL(16) or COMPLEX(16) Coarray

- Accessing another image's value of an ALLOCATABLE or POINTER component of a derived-type coarray

3.2.4 Static Security Analysis Feature (formerly Source Checker) Requires Intel® Inspector XE

The "Source Checker" feature, from compiler version 11.1, has been enhanced and renamed "Static Security Analysis". The compiler options to enable Static Security Analysis remain the same as in compiler version 11.1 (for example, `-diag-enable sc`), but the results are now written to a file that is interpreted by Intel® Inspector XE rather than being included in compiler diagnostics output.

3.2.5 New and Changed Directives (Update 6)

The following compiler directives are new or changed in Intel® Composer XE 2011 – please see the documentation for details:

- ATTRIBUTES VECTOR
- NOFUSION
- You can now specify a FIRSTPRIVATE clause in the PARALLEL directive
- You can now specify a FIRSTPRIVATE or LASTPRIVATE clause in the SIMD directive

3.2.6 OpenMP Changes (Update 6)

The following changes to OpenMP* support are in Intel® Composer XE 2011 :

- OpenMP 3.1 is supported
- TASKYIELD directive
- New clauses have been added to the ATOMIC directive
- You can now specify FINAL and MERGEABLE clauses in the TASK directive

3.2.7 Other Changes

- The ability to create a source listing file with identifier cross-reference has been added
- Guided auto-parallelism
- An option to use math library functions that are faster but return results with less precision or accuracy
- An option to use math library functions that return consistent results across different models and manufacturers of processors
- The ability to generate a build dependencies output file has been added

3.2.7.1 Change in Static Security Analysis Behavior

The `inspxe-runsc` command line utility that is distributed with Intel® Composer XE 2011 has been changed. This change only affects users who use Composer XE 2011 to perform Static Security Analysis (SSA). Those that do not use SSA and those that perform SSA without using this utility are unaffected. SSA is only available to users of Intel® Parallel Studio XE 2011 or Intel® C++ Studio XE 2011, so users who do not have those products are unaffected.

`inspxe-runsc` executes a **build specification,** a description of how an application is built. Usually build specification files are generated by observing a build as it executes and recoding the compilations and links that are performed. `inspxe-runsc` repeats these actions using the Intel compiler in SSA mode. SSA results are generated at the link step so a build specification that describes a build with more than one link step will generate more than one SSA result when `inspxe-runsc` is invoked.

The versions of `inspxe-runsc` included in Composer XE 2011 and Composer XE 2011 Update 1 generate all the SSA results in a single directory. In the multiple link case this violated the rule that all the SSA results for one and only one project must be created in the same directory. The updated

version of `inspxe-runsc` respects this rule by generating results for each link step in a separate directory. The name of that directory is formed from the name of the file being linked. Thus if a build specification describes a project that builds two executables, file1.out and file2.out, then earlier versions of `inspxe-runsc` would create two results, one for file1 and one for file2, say r000sc and r001sc, in the same directory. The new version of `inspxe-runsc` will also create two results, but the one for file1 will be created in "My Inspector XE results – file1/r000sc" and the one for file2 will be created in "My Inspector XE results – file2/r000sc". The directories containing the results are both created in the same parent directory.

`inspxe-runsc` has a command line switch, -result-dir (-r), that specifies where results are to be created. The meaning of this switch has changed. Previous this would name the directory where the result itself, say r000sc, would be created. Now it names the parent directory where the "My Inspector XE Results - name" directory or directories will be created. So the directory named in the –r switch is effectively two levels up from the results themselves.

The change to `inspxe-runsc` effectively moves the result directory, and user action is required to adapt to this change. Those using scripts that invoke `inspxe-runsc` with the –r switch must update their scripts to reflect the new interpretation of the –r switch argument described earlier. Users must move their old result files into the new directory so that SSA results produced by earlier versions of `inspxe-runsc` share the same directory as results produced by the new version of `inspxe-runsc`. Users that had been using `inspxe-runsc` with a build specification with only one link step should move their old results into a directory of the form "My Inspector XE results – name". If this is not done, then all the problems in the newly created result will appear to be "New". Users that had been using `inspxe-runsc` with a build specification with multiple link steps have been having various issues with SSA that will be resolved by using the new utility. Such users are best advised to copy the most recent into their old results into each of the new "My Inspector XE results – name" directories. This offers the best chance that some old problem state information will be correctly applied to new results when they are created in the future.

3.3 New and Changed Compiler Options

Please refer to the compiler documentation for details

3.3.1 New and Changed in Composer XE 2011 Update 6 (and later)

- -align [no]qcommons
- -assume [no]std_intent_in (Update 9)
- -f[no-]asynchronous-unwind-tables
- -axCORE-AVX-I
- -axCORE-AVX2
- -f[no-]fma
- -f[no-]merge-debug-strings
- -fopenmp
- -gdwarf-3
- -march=atom
- -march=core-avx-i
- -march=core-avx2
- -march=corei7-avx
- -march=corei7
- -march=Pentium-m
- -opt-mem-layout-trans[=n]

- -xCORE-AVX-I
- -xCORE-AVX2
- -xSSSE3_ATOM

3.3.2 New and Changed in Composer XE 2011 (through Update 5)

- -assume [no]fpe_summary
- -assume [no]old_ldout_format
- -coarray
- -coarray-num-images
- -fzero-initialized-in-bss
- -fimf-absolute-error
- -fimf-accuracy-bits
- -fimf-arch-consistency
- -fimf-max-error
- -fimf-precision
- -fvar-tracking
- -fvar-tracking-assignments
- -gen-dep
- -gen-depformat
- -guide
- -guide-data-trans
- -guide-file
- -guide-file-append
- -guide-opts
- -guide-par
- -guide-vec
- -list
- -list-line-len
- -list-page-len
- -opt-args-in-regs
- -par-runtime-control
- -prof-value-profiling
- -profile-functions
- -profile-loops-report
- -show=keyword
- -simd
- -sox=keyword
- -standard-semantics

For a list of deprecated compiler options, see the Compiler Options section of the documentation.

3.3.3 –assume [no]std_intent_in added (Update 9)

The `-assume [no]std_intent_in` option determines whether the compiler assumes that dummy arguments with the `INTENT(IN)` attribute are not modified across a call, in accordance with the Fortran standard. The default is `std_intent_in`, allowing the compiler to assume that `INTENT(IN)` arguments are not modified; `nostd_intent_in` tells the compiler not to make such assumptions. `-standard-semantics` implies `-assume std_intent_in`.

3.3.4 Additional Keywords for `-sox` option, default changed (Update 3)

The `-sox` option, which adds information to the object and executable file about compiler options used and procedure profiling information, has been enhanced to let the user request that the list of inlined functions be included and to let the user exclude information about procedure profiling.
The syntax for `-sox` is now:
`-[no-]sox`
`-sox=`*`keyword`*`[,`*`keyword`*`]`
Where *`keyword`* is one of `inline` or `profile`. If `-sox` is specified with no keywords, only the command line options are included – this is a change from previous releases. To maintain the previous behavior, use `-sox=profile`. Multiple `-sox` options may be specified on the command line – if so, they are interpreted in left-to-right order.

3.4 Other Changes and Notes

3.4.1 Optimization Reports Disabled by Default

As of version 11.1, the compiler no longer issues, by default, optimization report messages regarding vectorization, automatic parallelization and OpenMP threaded loops. If you wish to see these messages you must request them by specifying `-diag-enable vec`, `-diag-enable par` and/or `-diag-enable openmp`, or by using `-vec-report`, `-par-report` and/or `-openmp-report`.
Also, as of version 11.1, optimization report messages are sent to `stderr` and not `stdout`.

3.4.2 Establishing the Compiler Environment

The `compilervars.sh` script is used to establish the compiler environment.
The command takes the form:
`source <install-dir>/bin/compilervars.sh` *`argument`*
Where *`xxx`* is the package identifier and *`argument`* is either `ia32` or `intel64` as appropriate for the architecture you are building for. Establishing the compiler environment also establishes the environment for the Intel® Debugger, Intel® Performance Libraries and, if present, Intel® C++ Compiler.

3.4.3 OpenMP* Legacy Libraries Removed

The OpenMP "legacy" libraries have been removed in this release. Only the "compatibility" libraries are provided.

3.4.4 RANF Portability Function Is Now an Intrinsic

The RANF function in the portability library is a non-standard random number generator. As of the version 12.0 compiler, RANF is an intrinsic function with a new, higher-performance implementation. If your program has added USE IFPORT to provide access to RANF, no changes will be seen and you will get the older version. If your program does not have USE IFPORT, or you add INTRINSIC RANF, you will get the new version that returns a different sequence, for a given seed, than the older version. The portability subroutine SRAND is still used to set the seed for RANF. Intel recommends use of the standard intrinsic RANDOM_NUMBER, but RANF is provided for compatibility with applications already using it.

3.5 Known Issues

3.5.1 Coarray Issues

For a list of known issues with Fortran 2008 Coarray support, see Coarray Known Issues.

3.5.2 Allocatable Arrays and OpenMP* PRIVATE or FIRSTPRIVATE

The compiler may fail to properly initialize allocatable arrays named in OpenMP PRIVATE or FIRSTPRIVATE clauses. This issue will be corrected in a future update. If you encounter problems with this combination of features, try adding the option `-switch omp3_private` . This is a temporary workaround and should not be used on a permanent basis. The Intel issue ID for this problem is DPD200160978.

3.6 Fortran 2003 and Fortran 2008 Feature Summary

The Intel Fortran Compiler supports many features that are new in Fortran 2003. Additional Fortran 2003 features will appear in future versions. Fortran 2003 features supported by the current compiler include:

- The Fortran character set has been extended to contain the 8-bit ASCII characters ~ \ [ ] ` ^ { } | # @
- Names of length up to 63 characters
- Statements of up to 256 lines
- Square brackets [ ] are permitted to delimit array constructors instead of (/ /)
- Structure constructors with component names and default initialization
- Array constructors with type and character length specifications
- A named PARAMETER constant may be part of a complex constant
- Enumerators
- Allocatable components of derived types
- Allocatable scalar variables
- Deferred-length character entities
- PUBLIC types with PRIVATE components and PRIVATE types with PUBLIC components
- ERRMSG keyword for ALLOCATE and DEALLOCATE
- SOURCE= keyword for ALLOCATE
- Type extension
- CLASS declaration
- Polymorphic entities
- Inheritance association
- Deferred bindings and abstract types
- Type-bound procedures
- TYPE CONTAINS declaration
- ABSTRACT attribute
- DEFERRED attribute
- NON_OVERRIDABLE attribute
- GENERIC keyword for type-bound procedures
- FINAL subroutines
- ASYNCHRONOUS attribute and statement
- BIND(C) attribute and statement
- PROTECTED attribute and statement
- VALUE attribute and statement
- VOLATILE attribute and statement
- INTENT attribute for pointer objects

- Reallocation of allocatable variables on the left hand side of an assignment statement when the right hand side differs in shape or length (requires option `-assume realloc_lhs` if not deferred-length character)
- Bounds specification and bounds remapping on a pointer assignment
- ASSOCIATE construct
- SELECT TYPE construct
- In all I/O statements, the following numeric values can be of any kind: UNIT=, IOSTAT=
- NAMELIST I/O is permitted on an internal file
- Restrictions on entities in a NAMELIST group are relaxed
- Changes to how IEEE Infinity and NaN are represented in formatted input and output
- FLUSH statement
- WAIT statement
- ACCESS='STREAM' keyword for OPEN
- ASYNCHRONOUS keyword for OPEN and data transfer statements
- ID keyword for INQUIRE and data transfer statements
- POS keyword for data transfer statements
- PENDING keyword for INQUIRE
- The following OPEN numeric values can be of any kind: RECL=
- The following READ and WRITE numeric values can be of any kind: REC=, SIZE=
- The following INQUIRE numeric values can be of any kind: NEXTREC=, NUMBER=, RECL=, SIZE=
- Recursive I/O is allowed in the case where the new I/O being started is internal I/O that does not modify any internal file other than its own
- IEEE Infinities and NaNs are displayed by formatted output as specified by Fortran 2003
- BLANK, DECIMAL, DELIM, ENCODING, IOMSG, PAD, ROUND, SIGN, SIZE I/O keywords
- DC, DP, RD, RC, RN, RP, RU, RZ format edit descriptors
- In an I/O format, the comma after a P edit descriptor is optional when followed by a repeat specifier
- Rename of user-defined operators in USE
- INTRINSIC and NON_INTRINSIC keywords in USE
- IMPORT statement
- Allocatable dummy arguments
- Allocatable function results
- PROCEDURE declaration
- Procedure pointers
- ABSTRACT INTERFACE
- PASS and NOPASS attributes
- The COUNT_RATE argument to the SYSTEM_CLOCK intrinsic may be a REAL of any kind
- Execution of a STOP statement displays a warning if an IEEE floating point exception is signaling
- MAXLOC or MINLOC of a zero-sized array returns zero if the option `-assume noold_maxminloc` is specified.
- Type inquiry intrinsic functions
- COMMAND_ARGUMENT_COUNT intrinsic
- EXTENDS_TYPE_OF and SAME_TYPE_AS intrinsic functions
- GET_COMMAND intrinsic
- GET_COMMAND_ARGUMENT intrinsic
- GET_ENVIRONMENT_VARIABLE intrinsic

- IS_IOSTAT_END intrinsic
- IS_IOSTAT_EOR intrinsic
- MAX/MIN/MAXVAL/MINVAL/MAXLOC/MINLOC intrinsics allow CHARACTER arguments
- MOVE_ALLOC intrinsic
- NEW_LINE intrinsic
- SELECTED_CHAR_KIND intrinsic
- The following intrinsics take an optional KIND= argument: ACHAR, COUNT, IACHAR, ICHAR, INDEX, LBOUND, LEN, LEN_TRIM, MAXLOC, MINLOC, SCAN, SHAPE, SIZE, UBOUND, VERIFY
- ISO_C_BINDING intrinsic module
- IEEE_EXCEPTIONS, IEEE_ARITHMETIC and IEEE_FEATURES intrinsic modules
- ISO_FORTRAN_ENV intrinsic module

The following is a partial list of Fortran 2003 features that are unimplemented or are known not to work in this release.

- User-defined derived type I/O
- Parameterized derived types
- Default initialization of CLASS objects
- The keyword MODULE may be omitted in MODULE PROCEDURE
- Transformational intrinsics, such as MERGE and SPREAD, in initialization expressions

The Intel® Fortran Compiler also supports some features from the Fortran 2008 standard. Additional features will be supported in future releases. Fortran 2008 features supported by the current version include:

- Maximum array rank has been raised to 31 dimensions (Fortran 2008 specifies 15)
- Coarrays
- CODIMENSION attribute
- SYNC ALL statement
- SYNC IMAGES statement
- SYNC MEMORY statement
- CRITICAL and END CRITICAL statements
- LOCK and UNLOCK statements
- ERROR STOP statement
- ALLOCATE and DEALLOCATE may specify coarrays
- Intrinsic procedures IMAGE_INDEX, LCOBOUND, NUM_IMAGES, THIS_IMAGE, UCOBOUND
    - **Note:** ATOMIC_DEFINE and ATOMIC_REF are not supported in this version
- CONTIGUOUS attribute
- MOLD keyword in ALLOCATE
- DO CONCURRENT
- NEWUNIT keyword in OPEN
- G0 and G0.d format edit descriptor
- Unlimited format item repeat count specifier
- A CONTAINS section may be empty
- Intrinsic procedures BESSEL_J0, BESSEL_J1, BESSEL_JN, BESSEL_YN, BGE, BGT, BLE, BLT, DSHIFTL, DSHIFTR, ERF, ERFC, ERFC_SCALED, GAMMA, HYPOT, IALL, IANY, IPARITY, IS_CONTIGUOUS, LEADZ, LOG_GAMMA, MASKL, MASKR, MERGE_BITS, NORM2, PARITY, POPCNT, POPPAR, SHIFTA, SHIFTL, SHIFTR, STORAGE_SIZE, TRAILZ,

- Additions to intrinsic module ISO_FORTRAN_ENV: ATOMIC_INT_KIND, ATOMIC_LOGICAL_KIND, CHARACTER_KINDS, INTEGER_KINDS, INT8, INT16, INT32, INT64, LOCK_TYPE, LOGICAL_KINDS, REAL_KINDS, REAL32, REAL64, REAL128, STAT_LOCKED, STAT_LOCKED_OTHER_IMAGE, STAT_UNLOCKED
- An OPTIONAL dummy argument that does not have the ALLOCATABLE or POINTER attribute, and which corresponds to an actual argument that: has the ALLOCATABLE attribute and is not allocated, or has the POINTER attribute and is disassociated, or is a reference to the NULL() intrinsic function, is considered not present
- A dummy argument that is a procedure pointer may be associated with an actual argument that is a valid target for the dummy pointer, or is a reference to the intrinsic function NULL. If the actual argument is not a pointer, the dummy argument shall have the INTENT(IN) attribute.

4 Intel® Debugger (IDB)

The following notes refer to the Graphical User Interface (GUI) available for the Intel® Debugger (IDB) when running on IA-32 and Intel® 64 architecture systems. In this version, the `idb` command invokes the GUI – to get the command-line interface, use `idbc`.

4.1 Setting up the Java* Runtime Environment

The Intel® IDB Debugger graphical environment is a Java application and requires a Java Runtime Environment (JRE) to execute. The debugger will run with a version 5.0 (also called 1.5).
Install the JRE according to the JRE provider's instructions.
Finally you need to export the path to the JRE as follows:

```
export PATH=<path_to_JRE_bin_dir>:$PATH
```

4.2 Starting the Debugger

To start the debugger, first make sure that the compiler environment has been established as described at Establishing the Compiler Environment. Then use the command:

```
idb
```

or

```
idbc
```

as desired.
Once the GUI is started and you see the console window, you're ready to start the debugging session.
Note: Make sure, the executable you want to debug is built with debug info and is an executable file. Change permissions if required, e.g. `chmod +x <application_bin_file>`

4.3 Additional Documentation

Online help titled *Intel® Debugger Online Help* is accessible from the debugger graphical user interface as `Help > Help Contents`.
Context-sensitive help is also available in most debugger dialogs, indicated by a "?" button.

4.4 Debugger Features

4.4.1 Changes for Intel® Fortran Composer XE 2011 Update 6

4.4.1.1 Intel® Inspector XE 2011 Update 6 Supports "break into debug" with IDB

Intel® Inspector XE 2011 Update 6 now supports "break into debug" mode with the Intel® Composer XE 2011 Update 6 version of IDB. Refer to the Intel® Inspector XE 2011 Release Notes for more information.

4.4.1.2 Watchpoints now using processor debug registers (hardware based)

As of Intel® Fortran Composer XE 2011 Update 6 (IDB 12.1), watchpoint support is now entirely supported by using the processor debug registers. Their possible configurations are specified by the underlying processor architecture. The following architecture-related limitations are present; if possible, IDB will raise appropriate error messages to assist the user:

- Possible sizes of the watched memory areas are only 1, 2 , 4 or 8 (INTEL® 64 only) byte.
- The start address of the watched memory area has to be aligned with its size (e.g. it is not possible to watch 2 bytes starting with an odd address)
- There is only support for a maximum of four active/enabled watchpoints. Unused watchpoints can be disabled to free resources or to enable/create other ones.
- Only the following access modes are supported:
    - Write: trigger on write accesses
    - Any: trigger on either write or read accesses
    - Changed: trigger on write accesses that actually changed the value
- Watched memory areas must not overlap each other.
- Watchpoints are not scope related but tied to a process. As long as a process exists the watchpoints are active/enabled. Only if the process is terminated (e.g. rerun) the watchpoints will be disabled. They can be enabled again if the user wishes to do so.
- Using the debugger to access the watched memory area (e.g. assign a different value to a variable) bypasses the hardware detection. Hence watchpoints only trigger if the debuggee itself accessed the watched memory area.
- If the debuggee is running on a guest OS inside a virtual machine stepping over an instruction or code line might continue the process without stopping. Watchpoints are only guaranteed to work when the debuggee runs on real hardware.

4.4.2 Main Features of IDB

The debugger supports all features of the command line version of the Intel® IDB Debugger. Debugger functions can be called from within the debugger GUI or the GUI-command line. Please refer to the Known Limitations when using the graphical environment.

4.4.2.1 Threads Window

- Improved Data Sharing Dectection
- Support for OpenMP* 3.0
- Support for Linux* OS synchronization functions
- Improved data sharing detection analysis performance

4.4.2.2 Extended Breakpoints Feature

With this feature you can set breakpoints on routines in shared libraries which have not yet been loaded. The requested breakpoint will be realized whenever possible. You'll see unrealized breakpoints marked with a yellow triangle (not having an address, file and symbol name) in the GUI. On the command line those are marked as `<PENDING>`. Any ambiguity is directly resolved and you will get multiple realizations, e.g. requesting a breakpoint for an overloaded function. In the GUI, those are visualized as a tree with the requesting breakpoint as its node. On the command line the requesting breakpoint is marked as `<MULTIPLE>` and its realizations follow. Please note that for the command line this feature is only available in GDB mode.

4.4.2.3 Command `solib-search-path` now Implemented

The command line debugger idbc and the Command window of the GUI debugger now support the existing gdb command `solib-search-path` which is used to look up images or shared libraries when they have not been found in the usual places such as `$LD_LIBRARY_PATH`.
Please invoke the command line help to see the `solib-search-path` command usage:
(idb) help set solib-search-path
(idb) help show solib-search-path
or the abbreviated commands:
(idb) h set sol
(idb) h sho sol

4.4.2.4 New Command for Disassembly Style Display

The IDB debugger now provides two styles of disassembly views in the Assembler window or on the Command windows.
The new commands on the Command window are:
(idb) set disassembly-flavor [att|intel]
(idb) show disassembly-flavor
The commands can also be found by invoking the help:
(idb) help set
(idb) help show
In the GUI/Assembler window right-click 'Change Style' to switch between Intel and ATT style. ATT stands for AT&T style (also known as GNU style).

4.5 Known Issues and Changes

4.5.1 Coarray elements cannot be viewed.

The IDB Debugger cannot view coarray elements. Please refer to section 3.2.3.1 'How to Debug a Coarray Application' where a workaround is described.

4.5.2 Signals Dialog not working Signals Dialog not working

The Signals dialog accessible via the GUI dialog Debug / Signal Handling or the shortcut Ctrl+S is not working correctly. Please refer to the Intel® Debugger (IDB) Manual for use of the signals command line commands instead.

4.5.3 Resizing GUI

If the debugger GUI window is reduced in size, some windows may fully disappear. Enlarge the window and the hidden windows will appear again.

4.5.4 `$cdir, $cwd` Directories

`$cdir` is the compilation directory (if recorded). This is supported in that the directory is set; but `$cdir` is not itself supported as a symbol.
`$cwd` is the current working directory. Neither the semantics nor the symbol are supported.
The difference between `$cwd` and '.' is that `$cwd` tracks the current working directory as it changes during a debug session. '.' is immediately expanded to the current directory at the time an entry to the source path is added.

4.5.5 `info stack` Usage

The gdb mode debugger command `info stack` does not currently support negative frame counts the way gdb does, for the following command:

```
info stack [num]
```
A positive value of `num` prints the innermost `num` frames, a zero value prints all frames, and a negative value prints the innermost –`num` frames in reverse order.

4.5.6 `$stepg0` Default Value Changed

The debugger variable $stepg0 changed default to a value of 0. With the value "0" the debugger will step over code without debug information if you do a "step" command. Set the debugger variable to 1 to be compatible with previous debugger versions as follows:
(idb) set $stepg0 = 1

4.5.7 SIGTRAP error on some Linux* Systems

On some Linux distributions (e.g. Red Hat Enterprise Linux Server release 5.1 (Tikanga)) a SIGTRAP error may occur when the debugger stops at a breakpoint and you continue debugging. As a workaround you may define the SIGTRAP signal as follows on command line:

```
(idb) handle SIGTRAP nopass noprint nostop
SIGTRAP is used by the debugger.
SIGTRAP       No        No      No              Trace/breakpoint trap
(idb)
Caveat: With this workaround all SIGTRAP signals to the debuggee are
blocked.
```

4.5.8 idb GUI cannot be used to debug MPI processes

The idb GUI cannot be used to debug MPI processes. The command line interface (idbc) can be used for this purpose.

4.5.9 Thread Syncpoint Creation in GUI

While for plain code and data breakpoints the field “Location” is mandatory, thread syncpoints require both “Location” and “Thread Filter” to be specified. The latter specifies the threads to synchronize. Please note that for the other breakpoint types this field restricts the breakpoints created to the threads listed.

4.5.10 Stack Alignment for IA-32 Architecture

Due to changes in the default stack alignment for the IA-32 architecture, the usage of inferior calls (i.e. evaluation of expressions that cause execution of debuggee code) might fail. This can cause as well crashes of the debuggee and therefore a restart of the debug session. If you need to use this feature, make sure to compile your code with 4 byte stack alignment by proper usage of the `–falign-stack=<mode>` option.

4.5.11 GNOME Environment Issues

With GNOME 2.28, debugger menu icons may not being displayed by default. To get the menu icons back, you need to go to the “System->Preferences->Appearance, Interface” tab and enable, "Show icons in menus". If there is not “Interface” tab available, you can change this with the corresponding `GConf` keys in console as follows:

```
gconftool-2 --type boolean --set /desktop/gnome/interface/buttons_have_icons true
gconftool-2 --type boolean --set /desktop/gnome/interface/menus_have_icons true
```

4.5.12 Accessing Online-Help

On systems where the Online-Help is not accessible from the IDB Debugger GUI Help menu, you can access the web-based debugger documentation from http://intel.ly/ng91lO

4.5.13 Debugger crashes if $HOME not set on calling shell

The debugger will end with a "Segmentation fault" if no $HOME environment variable is set on the shell the debugger is started from.

4.5.14 Command line parameter –parallel not supported

The debugger command line parameter –parallel is not supported on the shell command prompt nor on the Console Window of the Debugger GUI.

4.5.15 Command line parameter –idb and -dbx not supported

The debugger command line parameters –idb and -debx are not supported in conjunction with the debugger GUI.

4.5.16 Core File Debugging

To be able to debug core files you need to start the debugger (command line debugger idbc or GUI debugger idb) with commandline options as follows:

```
idb|idbc <executable> <corefile>
```

<or>

```
idb|idbc <executable> –core <corefile>
```

Once started with a core file, the debugger is not able to debug a live process e.g. attaching or creating a new process. Also when debugging a live process, a core file cannot be debugged.

4.5.17 Thread Data Sharing Filters may not work correctly

Setting Thread Data Sharing Filters may lead to unexpected behavior of the debugger. It may happen that threads will not continue after a data sharing detection and the debugger may exit with a SIGSEGV. If you encounter issues related to Data Sharing Detection with filters enabled, disable all filters in the 'Thread Data Sharing Filters' window context menu.

4.5.18 Default .gdbinit script on Pardus systems may cause the Debugger crash

If you encounter a debugger crash when starting idbc or idb, you may add the option `–nx` to bypass the default .gdbinit script.

4.5.19 No thread info available on Pardus systems

Due to an issue with the default libthread_db.so library on Pardus systems, the debugger cannot detect thread info when debugging multithreaded applications.

5 Intel® Math Kernel Library

This section summarizes changes, new features and late-breaking news about this version of the Intel® Math Kernel Library (Intel® MKL). For information on bug fixes, see http://intel.ly/neQIw2

5.1 What's New in Intel® MKL 10.3 Update 9

- LAPACK: Improved [C/Z]GEEV performance for very small sizes (~10 x 10)
- FFTs: Threaded the real in-place 1D FFTs for a significant increase in performance
- FFTs: Introduced new algorithms for improved scalability of power-of-2 double-precision complex 1D FFTs on Intel® Xeon® Processor E5 series systems running 32-bit operating systems
- Random number Generators: added support for a non-deterministic random number generator based on the RdRand instruction and supporting hardware available in future processors based on the Intel® code name "Ivy Bridge" microarchitecture

- Vector Math Functions: improved performance of the Erf() and Pow3o2() functions on Intel® Core™ processors
- Data Fitting: improved performance of routines for spline-based evaluation, differentiation, and integration on Intel® Xeon® 5600 and 7500 series and Intel® Core™ i7-2600 series processors

5.2 What's New in Intel® MKL 10.3 Update 8

- Data Fitting component: Added a set of new data fitting functions covering one-dimensional algorithms for vector spline construction, cell or bin search, and evaluation, differentiation, and integration of the spline interpolants. Includes support for:
  - Linear, quadratic, cubic, step-wise const, and user-defined splines
  - Cell search with configuration parameters for optimal performance
  - User-defined interpolation and extrapolation
  - Vector-valued functions
  - Column- and row-major storage formats
- Sparse BLAS: Improved compressed sparse row matrix-vector multiply (?CSRMV) performance for very sparse matrices on high core counts supporting Intel Advanced Vector Extensions (AVX)
- FFTs: Improved the performance of the 1D double precision FFTs on systems supporting Intel AVX
- Statistics functions: Improved the performance and scalability for computing the Variance-Covariance and Correlation matrices (FAST method) on Intel® Core processors
- Bug fixes

5.3 What's New in Intel® MKL 10.3 Update 7

- BLAS: Improved DSYRK/SSYRK threaded performance for small output matrices and large outer products (i.e., rectangular input matrices), on all recent Intel® Xeon® processors
- BLAS: Improved ?GEMM performance for small problems (<10) where beta =1 on all recent Intel Xeon processors
- BLAS: Improved DSCAL performance for small problems and for cases where INCX=1 on 32-bit programs running on Intel Xeon processors 5500, 5600, and 7500 series
- BLAS-like extensions: Improved threading and cache utilization of in-place transposition of square matrices
- PARDISO: Introduced an independent threading control for PARDISO; use MKL_DOMAIN_PARDISO with the `mkl_domain_set_num_threads()` function
- Poisson Library: Added support for 2D and 3D periodic boundary conditions
- Included the Link Line Advisor in the documentation directory
- Added a command line link tool for use with scripting tools such as libtool
- Changed the names of constants used to specify the domain in the `mkl_domain_set_num_threads()` function (e.g., MKL_BLAS has become MKL_DOMAIN_BLAS); the old names still exist with the exception of MKL_PARDISO
- Bug fixes

5.4 What's New in Intel® MKL 10.3 Update 6

- Sparse BLAS: Added a new option to the mkl_?csrbsr converter function allowing detection and removal of zero elements when converting from the BSR format to the CSR format
- Bug fixes

5.5 What's New in Intel® MKL 10.3 Update 5

- BLAS: Improved performance: {S,C,Z}TRSM for processors with Intel® Advanced Vector Extensions (Intel® AVX); {S,D}GEM2VU for processors with Intel AVX as well as the Intel® Core™ i7 processor and the Intel® Xeon® processor 5500 series
- BLAS: Improved scaling: ?TRMV for large matrices on all architectures; DGEMM for odd numbers of threads on Intel® Xeon® processor 5400 series
- LAPACK: Included LAPACK 3.3.1 extensions and the respective LAPACKE interfaces
- LAPACK: Improved the performance of ?SYGST and ?HEGST used in generalized eigenvalue problems
- LAPACK: Improved the performance of the inverse of an LU factored matrix (?GETRI)
- PARDISO: Added transpose and conjugate transpose solve capability (ATx=b and AHx=b); facilitates compressed sparse column (CSC) format support
- PARDISO: Improve out-of-core PARDISO performance when the memory requirements slightly exceed available memory using MKL_PARDISO_OOC_MAX_SWAP_SIZE environment variable and in-core PARDISO
- Optimization Solvers: Added Inf and NaN checks in the RCI Trust-Region solvers
- FFTs: Improved the performance of 3D FFTs on small cubes from 2x2x2 to 10x10x10 for all supported precisions and types on all Intel® processors supporting Intel® SSE3 and later
- FFT examples: Re-designed example programs to cover common use cases for Intel MKL DFTI and FFTW
- VSL: Improved the performance of the single precision MT19937 and MT2203 basic random number generators on the Intel® Core™ i7-2600 processor on 64-bit operating systems
- VSL: Improved the performance of the integer version of the SOBOL quasi-random number generator on the Intel® Core™ i7-2600 processor and Intel® Xeon® processor 5400 series
- Bug fixes

5.6 What's New in Intel® MKL 10.3 Update 4

- BLAS: Improved DTRMM performance on Intel® Xeon® processors 5400 and later
- BLAS: Improved DTRSM performance on all 64-bit enabled processors, especially processors with Intel® Advanced Vector Extensions (Intel® AVX)
- LAPACK: Incorporated bug fixes from the LAPACK 3.3.1 release
- OOC PARDISO: Improved the estimate of the amount of memory needed in out-of-core operation
- FFT: Improved 1D real FFT scaling through improved threading
- FFT: Updated C and Fortran FFT examples to use the new single dynamic library linking model
- VML: Improved performance of the single precision Enhanced Performance version of the real Hypot and complex Abs functions and of the complex Arg, Div, Mul, MulByConj functions for all accuracy modes on Intel® Xeon® processors 5600 and 7500 series, and the Intel® Core™ i7-2600 processor
- Service functions: Improvements and additions to the Intel MKL service functions the online release notes at http://intel.ly/pkUQXI for more information)
- Bug fixes

5.7 What's New in Intel® MKL 10.3 Update 3

- BLAS: Improved multi-threaded performance of DSYRK, DTRSM, and DGEMM on Intel® Xeon® processor 5400 series running 32-bit Windows*

- LAPACK: Implemented LAPACK 3.3 from netlib including Cosine-Sine decomposition, improved linear equations solvers for symmetric and Hermitian matrices and auxiliary functions
- PARDISO: 0-based permutation vectors are now allowed at input
- PARDISO: Documentation for the pardisoinit() routine
- PARDISO: Improved performance of serial PARDISO with multiple right-hand sides (RHS)
- PARDISO: Independent control for parallelism in the solve step for improved performance on small matrices—see description of iparm(25)
- PARDISO: Reduced backward substitution—allows partial solution computation for a full RHS—see description of iparm(31)
- FFT: Implemented Real FFT transforms for up to 3 to 7 dimensions
- FFT: Parallelized multi-dimensional complex transforms using split-complex data represented as two real arrays
- Cluster FFTs: Extended FORTRAN 90 interface to real-to-complex transforms and included new examples
- VML: Added new complex Pack/Unpack functions and real Gamma/LGamma functions
- VML: Improved performance on Intel® Xeon® processor 5600 series and processors supporting Intel® Advanced Vector Extensions (Intel® AVX) for the following: all functions when operating on short vectors (<100), all functions when operating on unaligned input vectors, the sPow2o3 function, and the enhanced performance (EP) version of complex Add and Sub.
- VSL: Functions for saving/restoring random number generator (RNG) streams to/from memory
- VSL: Added new UniformBits32 and UniformBits64 functions
- VSL: Extended the number of unique streams supported by MT2203 BRNG from 1024 to 6024
- Bug fixes

Note: The GMP Arithmetic Functions in Intel MKL will be removed in a future version of Intel MKL.

5.8 What's New in Intel® MKL 10.3 Update 2

- BLAS: Improved performance of transposition functions on the Intel® Xeon® processor 5600 series
- BLAS: Added examples for transposition routines
- FFT: Added Fortran examples showing how to reduce application footprint by linking only functions with the desired precision
- FFT: Added check for stride consistency on in-place real transforms with CCE storage
- FFT: Expanded threading to new cases for multi-dimensional transforms
- VSL: Improved performance of Multivariate Gaussian random number generator for single- and double-precision on 4-core Intel® Xeon® processors 5500 series
- VML: Improved performance of in-place operation of Add, Mul, and Sub functions on the Intel® Xeon® processor 5500 series
- Bug fixes

5.9 What's New in Intel® MKL 10.3 Update 1

- PARDISO/DSS: Added true F90 overloaded API (see the Intel® MKL reference manual for more information)
- PARDISO: Improved the statistical reporting to be more reader friendly
- Sparse BLAS: Improved performance of ?BSRMM functions on the latest Intel& processors
- FFTs: Support for negative strides

- FFT examples: Added examples for split-complex FFTs in C and Fortran using both the DFTI and FFTW3 interfaces
- VML: Improved performance of real in-place Add/Sub/Mul/Sqr functions on systems supporting SSE2 and SSE3
- Poisson Library: Changed the default behavior of the Poisson library functions from sequential to threaded operation
- Bug fixes

5.10 What's New in Intel® MKL 10.3

- BLAS
  - New functions for computing 2 matrix-vector products at once: [D/S]GEM2VU, [Z/C]GEM2VC
  - New functions for computing mixed precision general matrix-vector products: [DZ/SC]GEMV
  - New function for computing the sum of two scaled vectors: *AXPBY
  - Intel® AVX optimizations in key functions: SMP LINPACK, level 3 BLAS, DDOT, DAXPY
- LAPACK
  - New C interfaces for LAPACK supporting row-major ordering
  - Integrated Netlib LAPACK 3.2.2 including one new computational routine (*GEQRFP) and two new auxiliary routines (*GEQR2P and *LARFGP) and the earlier LAPACK 3.2.1 update
  - Intel® AVX optimizations in key functions: DGETRF, DPOTRF, DGEQRF
- PARDISO
  - Improved performance of factor and solve steps in multi-core environments
  - Introduced the ability to solve for sparse right-hand sides and perform partial solves—produces partial solution vector
  - Improved performance of the out-of-core (OOC) factorization step
  - Support for zero-based (C-style) array indexing
  - Zeros on the diagonal of the matrix are no longer required in sparse data structures for symmetric matrices
  - New ILP64 PARDISO interface allows the use of both LP64 and ILP64 versions when linked to the LP64 libraries
  - The memory required for storing files on the disk in OOC mode can now be estimated just after reordering
- Sparse BLAS
  - Format conversion functions now support all data types (single and double precision for real and complex data) and can return sorted or unsorted arrays
- FFTs
  - Intel AVX optimizations in all 1D/2D/3D FFTs
  - Improved performance of 2D and 3D mixed-radix FFTs for single and double precision data for all systems supporting the SSE4.2 instruction set
  - Support for split-complex data represented as two real arrays introduced for 2D/3D FFTs
  - Support for 1D complex-to-complex transforms of large prime lengths
- VML

    - A new function for computing (ax+b)/(cy+d) where a, b, c, and d are scalars, and x and y are real vectors: v[s/d]LinearFrac()
    - Intel AVX optimizations for real functions
    - A new mode for setting denormals to zero, overflow support for complex vectors, and for every VML function a new function with an additional parameter for setting the accuracy mode
- VSL
    - A set of new Summary Statistics functions was added covering basic statistics, covariance and correlation, pooled, group, partial, and robust covariance/correlation, quantiles and streaming quantiles, outliers detection algorithm, and missing values support
        - Performance optimized algorithms: MI algorithm for support of missing values, TBS algorithm for computation of robust covariance, BACON algorithm for detection of outliers, ZW algorithm for computation of quantiles (streaming data case), and 1PASS algorithm for computation of pooled covariance
    - Improved performance of SFMT19937 Basic Random Number Generator (BRNG)
    - Intel® AVX optimizations: MT19937 and MT2203 BRNGs
- Added runtime dispatching dynamic libraries allowing link to a single interface library which loads dependent libraries dynamically at runtime depending on runtime CPU detection and/or library function calls
- The custom dynamic libraries builder now uses the runtime dispatching dynamic libraries on the Linux* and Mac OS* X operating systems
- A new directory structure has been established to simplify integration of Intel MKL with the Intel® Parallel Studio XE family of products and directories formerly designated as "em64t" are now designated by the "intel64" tag
- The sparse solver functionality has been fully integrated into the core Intel MKL libraries and the libraries with "solver" in the filename have been removed from the product

5.11 Known Issues

A full list of the known limitations of this release can be found in the Knowledge Base for the Intel® MKL at http://intel.ly/ptEfAP

5.12 Notices

The following changes are planned for future versions of Intel MKL. Please contact Technical Support if you have concerns:

- Content in the libraries containing solver in the filenames will be moved to the core library in a future version of Intel MKL. These solver libraries will then be removed.
- The Intel MKL GNU Multiple Precision* (GMP) function interfaces will be removed in a future library release.
- The timing function `mkl_set_cpu_frequency()` is deprecated. Please use `mkl_get_max_cpu_frequency()`, `mkl_get_clocks_frequency()`, and `mkl_get_cpu_frequency()` as described in the Intel® MKL Reference Manual.
- The MKL_PARDISO constant defined to specify the PARDISO domain should no longer be used with the mkl_domain_set_num_threads() function; please use MKL_DOMAIN_PARDISO instead.
- Convolution and Correlation routines will not be backward compatible with 10.2 update 3 in a future release.

5.13 Attributions

As referenced in the End User License Agreement, attribution requires, at a minimum, prominently displaying the full Intel product name (e.g. "Intel® Math Kernel Library") and providing a link/URL to the Intel® MKL homepage (www.intel.com/software/products/mkl) in both the product documentation and website.

The original versions of the BLAS from which that part of Intel® MKL was derived can be obtained from http://www.netlib.org/blas/index.html.

The original versions of LAPACK from which that part of Intel® MKL was derived can be obtained from http://www.netlib.org/lapack/index.html. The authors of LAPACK are E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, and D. Sorensen. Our FORTRAN 90/95 interfaces to LAPACK are similar to those in the LAPACK95 package at http://www.netlib.org/lapack95/index.html. All interfaces are provided for pure procedures.

The original versions of ScaLAPACK from which that part of Intel® MKL was derived can be obtained from http://www.netlib.org/scalapack/index.html. The authors of ScaLAPACK are L. S. Blackford, J. Choi, A. Cleary, E. D'Azevedo, J. Demmel, I. Dhillon, J. Dongarra, S. Hammarling, G. Henry, A. Petitet, K. Stanley, D. Walker, and R. C. Whaley.

PARDISO in Intel® MKL is compliant with the 3.2 release of PARDISO that is freely distributed by the University of Basel. It can be obtained at http://www.pardiso-project.org.

Some FFT functions in this release of Intel® MKL have been generated by the SPIRAL software generation system (http://www.spiral.net/) under license from Carnegie Mellon University. The Authors of SPIRAL are Markus Puschel, Jose Moura, Jeremy Johnson, David Padua, Manuela Veloso, Bryan Singer, Jianxin Xiong, Franz Franchetti, Aca Gacic, Yevgen Voronenko, Kang Chen, Robert W. Johnson, and Nick Rizzolo.

6 Disclaimer and Legal Information

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL(R) 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. UNLESS OTHERWISE AGREED IN WRITING BY INTEL, THE INTEL PRODUCTS ARE NOT DESIGNED NOR INTENDED FOR ANY APPLICATION IN WHICH THE FAILURE OF THE INTEL PRODUCT COULD CREATE A SITUATION WHERE PERSONAL INJURY OR DEATH MAY OCCUR.

Intel may make changes to specifications and product descriptions at any time, without notice. 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 information here is subject to change without notice. Do not finalize a design with this information.

The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or go to:
http://www.intel.com/design/literature.htm
Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families. Go to:
http://www.intel.com/products/processor_number/ for details.
Celeron, Centrino, Intel, Intel logo, Intel386, Intel486, Intel Atom, Intel Core, Itanium, MMX, Pentium, VTune, and Xeon are trademarks of Intel Corporation in the U.S. and other countries.
* Other names and brands may be claimed as the property of others.
Copyright © 2012 Intel Corporation. All Rights Reserved.