

TAU Install Guide

TAU Install Guide

Updated October 20th 2008, for use with Version 2.18.3.

Copyright © 1997-2008 Department of Computer and Information Science, University of Oregon Advanced Computing Laboratory, LANL, NM Research Centre Julich, ZAM, Germany

Permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of University of Oregon (UO) Research Centre Julich, (ZAM) and Los Alamos National Laboratory (LANL) not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission. The University of Oregon, ZAM and LANL make no representations about the suitability of this software for any purpose. It is provided "as is" without express or implied warranty.

UO, ZAM AND LANL DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL THE UNIVERSITY OF OREGON, ZAM OR LANL BE LIABLE FOR ANY SPECIAL, INDIRECT OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

TAU can be found on the web at: <http://www.cs.uoregon.edu/research/tau>

Table of Contents

1. Installation	1
1.1. Installing TAU	1
1.1.1. Know what options you will need	1
1.1.2. Common configuration options	2
1.1.3. Configuring with external packages	3
1.1.4. More configuration options	4
1.1.5. tau_setup	8
1.1.6. installtau script	8
1.1.7. upgradetau	9
1.1.8. tau_validate	9
1.2. Platforms Supported	9
1.3. Software Requirements	10

Chapter 1. Installation

TAU (Tuning and Analysis Utilities) is a portable profiling and tracing toolkit for performance analysis of parallel programs written in Fortran, C++, C, Java and Python. The model that TAU uses to profile parallel, multi-threaded programs maintains performance data for each thread, context, and node in use by an application. The profiling instrumentation needed to implement the model captures data for functions, methods, basic blocks, and statement execution at these levels. All C++ language features are supported in the TAU profiling instrumentation including templates and namespaces, which is available through an API at the library or application level. The API also provides selection of profiling groups for organizing and controlling instrumentation. The instrumentation can be inserted in the source code using an automatic instrumentor tool based on the Program Database Toolkit (PDT), dynamically using DyninstAPI, at runtime in the Java virtual machine, or manually using the instrumentation API. TAU's profile visualization tool, paraprof, provides graphical displays of all the performance analysis results, in aggregate and single node/context/thread forms. The user can quickly identify sources of performance bottlenecks in the application using the graphical interface. In addition, TAU can generate event traces that can be displayed with the Vampir or Paraver trace visualization tools. This chapter discusses installation of the TAU portable profiling package.

1.1. Installing TAU

1.1.1. Know what options you will need

Each TAU configuration results in a single Makefile. These Makefiles denote the configuration that produced it and is used by the user to select the TAU libraries/scripts associated with its configuration. (These makefiles are named after the configuration options, ie. TAU configured with MPI, PDT, PGI compilers and the '-nocomm' option is named: `Makefile.tau-nocomm-mpi-pdt-pgi`). On most machines several configuration of TAU will need to be built in order to take full advantage of the many features of TAU. This section should help you decide on the smallest set of configuration you will need to build.

The options used to configure TAU can be grouped into two categories:

- External packages: TAU will use these when instrumenting or measuring an application. *Configuring with these options does not force the user to use these packages*, ie: configuring with PDT does not force the user to use source code based instrumentation (they can use compiler based instrumentation instead). Similarly configuring with PAPI does not forces the user to select any PAPI counters when profiling.



Note

The only exception is configuring with the epilog (scalasca) tracing package. This will replace the TAU tracer with the eplilog one, a single configuration cannot use both tracers.

For this reason it is recommend that you *configure with every external packages that the user might be interested in using*, letting them choose which packages to enable when they go to instrument or measure their application.

- Compiler and MPI options: these control the behavior of TAU when it compiles the instrumented application. TAU provides compiler wrapper scripts, these options control which compiler TAU will wrap, *These options are determinative: select only options that are compatible*. For example, when configuring with MPI use a version of MPI compatible with the compiler you select.

Since multiple compiler/MPI libraries cannot be specified for a single configuration, *each set of compiler/MPI libraries that you want to use with TAU need to be configured separately.*



Note

Configurations with different compilers are given separate Makefiles automatically, however configurations with different MPI implementations are not. Use the `-tag=` option to distinguish between different MPIs, ie: `-tag=mvapich` or `-tag=openmpi`.

The `configure` shell script attempts to guess correct values for various system-dependent variables used during compilation (compilers and system architecture), other options need to be specified on the command line.

The following are the most important command-line options are available to configure:

1.1.2. Common configuration options

1.1.2.1. Select compiler

TAU defaults to using any compilers found in the environment. To use a specific compiler use these options:

- `-c++=<C++ compiler>`

Specifies the name of the C++ compiler. Supported C++ compilers include KCC (from KAI/Intel), CC (SGI, Sun), g++ (from GNU), FCC (from Fujitsu), xlc (from IBM), guidec++ (from KAI/Intel), cxx (Tru64) and aCC (from HP), c++ (from Apple), icpc and ecpc (from Intel) and pgCC (from PGI).

- `-cc=<C Compiler>`

Specifies the name of the C compiler. Supported C compilers include cc, gcc (from GNU), pgcc (from PGI), fcc (from Fujitsu), xlc (from IBM), and KCC (from KAI/ Intel), icc and ecc (from Intel).

- `-fortran=<Fortran Compiler>`

Specifies the name of the Fortran90 compiler. Valid options are: gnu, sgi, ibm, ibm64, intel, cray, pgi, absoft, fujitsu, sun, kai, nec, hitachi, compaq, nagware, g95 and hp.

1.1.2.2. Using MPI

To profile program that use MPI configure TAU with these options:

- `-mpi`

With this options TAU will try to guess the location of the MPI libraries if the `mpirun` command is in your path. This does not always work in which case use these more detailed options:

- `-mpiinc=<dir>`

Specifies the directory where MPI header files reside (such as `mpi.h` and `mpif.h`). This option also

generates the TAU MPI wrapper library that instruments MPI routines using the MPI Profiling Interface. See the `examples/NPB2.3/config/make.def` file for its usage with Fortran and MPI programs. MPI [<http://www-unix.mcs.anl.gov/mpi/>]

- `-mpilib=<dir>`

Specifies the directory where MPI library files reside. This option should be used in conjunction with the `-mpiinc=<dir>` option to generate the TAU MPI wrapper library.

- `-mpilibrary=<lib>`

Specifies the use of a different MPI library. By default, TAU uses `-lmpi` or `-lmpich` as the MPI library. This option allows the user to specify another library. e.g., `-mpilibrary=-lmpi_r` for specifying a thread-safe MPI library.

1.1.2.3. OpenMP

TO profile programs using openmp use *both* of these options:

- `-openmp`

Specifies OpenMP as the threads package to be used. Open MPI [<http://www.open-mpi.org/>]

- `-opari`

The use of Opari source-to-source instrumentor in conjunction with TAU exposes OpenMP events for instrumentation. See `examples/opari` directory. OPARI [<http://www.fz-juelich.de/zam/kojak/opari/>]

1.1.3. Configuring with external packages

- `-pdt=<directory>`

Specifies the location of the installed PDT (Program Database Toolkit) root directory. PDT is used to build `tau_instrumentor`, a C++, C and F90 instrumentation program that automatically inserts TAU annotations in the source code PDT [<http://www.cs.uoregon.edu/research/pdt/>]. If PDT is configured with a subdirectory option (`-compdir=<opt>`) then TAU can be configured with the same option by specifying

- `-papi=<directory>`

Specifies the location of the installed PAPI (Performance Data Standard and API) root directory. PCL provides a common interface to access hardware performance counters and timers on modern microprocessors. Most modern CPUs provide on-chip hardware performance counters that can record several events such as the number of instructions issued, floating point operations performed, the number of primary and secondary data and instruction cache misses, etc. To measure floating point instructions, set the environment variable `TAU_METRICS` to `PAPI_FP_INS` (for example). This option (by default) specifies the use of hardware performance counters for profiling (instead of time). PAPI [<http://icl.cs.utk.edu/papi/>]

- `-epilog=<dir>`

Specifies the directory where the EPILOG tracing package EPILOG

[<http://www.fz-juelich.de/zam/kojak/epilog/>] is installed. This option should be used in conjunction with the `-TRACE` option to generate binary EPILOG traces (instead of binary TAU traces). EPILOG traces can then be used with other tools such as EXPERT. EPILOG comes with its own implementation of the MPI wrapper library and the POMP library used with Opari. Using option overrides TAU's libraries for MPI, and OpenMP.

- `-otf=<directory>`

Specifies the location of the OTF trace libraries generation package. TAU's binary traces can be converted to the OTF format using `tau2otf`, a tool that links with the OTF library.

- `-vtf=<directory>`

Specifies the location of the VTF3 trace generation package. TAU's binary traces can be converted to the VTF3 format using `tau2vtf`, a tool that links with the VTF3 library. The VTF3 format is read by Intel trace analyzer, formerly known as vampir, a commercial trace visualization tool developed by TU. Dresden, Germany.

1.1.4. More configuration options

- `-PROFILEPHASE`

This option generates phase based profiles. It requires special instrumentation to mark phases in an application (I/O, computation, etc.). Phases can be static or dynamic (different phases for each loop iteration, for instance). See `examples/phase/README` for further information.

- `-prefix=<directory>`

Specifies the destination directory where the header, library and binary files are copied. By default, these are copied to subdirectories `<arch>/bin` and `<arch>/lib` in the TAU root directory.

- `-arch=<architecture>`

Specifies the architecture. If the user does not specify this option, configure determines the architecture. For IBM BGL, the user should specify `bgl` as the architecture. For SGI, the user can specify either of `sgi32`, `sgin32` or `sgi64` for 32, `n32` or 64 bit compilation modes respectively. The files are installed in the `<architecture>/bin` and `<architecture>/lib` directories.

- `-pdt_c++=<C++ Compiler>`

Specifies a different C++ compiler for PDT (`tau_instrumentor`). This is typically used when the library is compiled with a C++ compiler (specified with `-c++`) and the `tau_instrumentor` is compiled with a different `<pdt_c++>` compiler. For e.g.,

```
-c++=pgCC -cc=pgcc -pdt_c++=KCC -openmp ...
```

uses PGI's OpenMP compilers for TAU's library and KCC for `tau_instrumentor`.

```
-arch=bgl -pdt=/usr/pdtoolkit-3.4 -pdt_c++=xlc -mpi
```

uses PDT, MPI for IBM BG/L and specifies the use of the front-end xlc compiler for building `tau_instrumentor`.

- `-tag=<Unique Name>`

Specifies a tag in the name of the stub Makefile and TAU makefiles to uniquely identify the installation. This is useful when more than one MPI library may be used with different versions of compilers. e.g.,

```
% configure -c++=icpc -cc=icc -tag=intel71-vmi \
             -mpiinc=/vmi2/mpich/include
```

- `-pthread`
Specifies pthread as the thread package to be used. In the default mode, no thread package is used.
- `-tulipthread=<directory> -smarts`
Specifies SMARTS (Shared Memory Asynchronous Runtime System) as the threads package to be used. `<directory>` gives the location of the SMARTS root directory. Smarts [<http://www.cs.uoregon.edu/research/tau/users/smarts.php>]
- `-opari_region`
Report performance data for only OpenMP regions and not constructs. By default, both regions and constructs are profiled with Opari.
- `-opari_construct`
Report performance data for only OpenMP constructs and not Regions. By default, both regions and constructs are profiled with Opari.
- `-pdtarch=<architecture>`
Specifies the architecture used to build pdt, default the tau architecture.
- `-papithreads` - Same as papi, except uses threads to highlight how hardware performance counters may be used in a multi-threaded application. When it is used with PAPI, TAU should be configured with `-papi=<dir> -pthread` autoinstrument Shows the use of Program Database Toolkit (PDT) for automating the insertion of TAU macros in the source code. It requires configuring TAU with the `-pdt=<dir>` option. The Makefile is modified to illustrate the use of a source to source translator (`tau_instrumentor`).
- `-jdk=<directory>`
Specifies the location of the installed Java 2 Development Kit (JDK1.2+) root directory. TAU can profile or trace Java applications without any modifications to the source code, byte-code or the Java virtual machine. See README.JAVA on instructions on using TAU with Java 2 applications. This option should only be used for configuring TAU to use JVMPI for profiling and tracing of Java applications. It should not be used for configuring paraprof, which uses Java from the user's path.
- `-dyninst=<dir>`
Specifies the directory where the DynInst dynamic instrumentation package is installed. Using DynInst, a user can invoke `tau_run` to instrument an executable program at runtime or prior to execution by rewriting it. DyninstAPI [<http://www.dyninst.org/>] PARA-DYN [<http://www.paradyn.org/>].
- `-vampirtrace=<directory>`
Specifies the location of the Vampir Trace package. With this option TAU will generate traces in Open Trace Format (OTF). More information at Technische Universitat Dresden [ht-

[tp://www.tu-dresden.de/zih/vampirtrace\]](http://www.tu-dresden.de/zih/vampirtrace)

- `-slog2=<directory>`

Specifies the location of the SLOG2 SDK trace generation package. TAU's binary traces can be converted to the SLOG2 format using `tau2slog2`, a tool that uses the SLOG2 SDK. The SLOG2 format is read by the Jumpshot4 trace visualization software, a freely available trace visualizer from Argonne National Laboratories. [Ref: <http://www-unix.mcs.anl.gov/perfvis/download/index.htm#slog2sdk>]

- `-slog2`

Specifies the use of the SLOG2 trace generation package and the Jumpshot trace visualizer that is bundled with TAU. Jumpshot v4 and SLOG2 v1.2.5delta are included in the TAU distribution. When the `-slog2` flag is specified, `tau2slog2` and `jumpshot` tools are copied to the `<tau>/<arch>/<bin>` directory. It is important to have a working `javac` and Java (preferably v1.4+) in your path. On Linux systems, where `/usr/bin/java` may be a place holder, you'll need to modify your path accordingly.

- `-mpilib= <lib>`

Specifies the use of a different MPI library. By default, TAU uses `-lmpi` or `-lmpich` as the MPI library. This option allows the user to specify another library. e.g., `-mpilib=-lmpi_r` for specifying a thread-safe MPI library.

- `-shmемinc=<dir>`

Specifies the directory where `shmem.h` resides. Specifies the use of the TAU SHMEM interface.

- `-shmemlib=<dir>`

Specifies the directory where `libsma.a` resides. Specifies the use of the TAU SHMEM interface.

- `-shmemlibrary=<lib>`

By default, TAU uses `-lsma` as the `shmem/pshmem` library. This option allows the user to specify a different `shmem` library.

- `-nocomm`

Allows the user to turn off tracking of messages (synchronous/asynchronous) in TAU's MPI wrapper interposition library. Entry and exit events for MPI routines are still tracked. Affects both profiling and tracing.

- `-epiloglib=<dir>`

Specifies the directory of where the Epilog library is to be found. Ex: if directory structure is: `/usr/local/epilog/fe/lib/` let the install options be: `-epilog=/usr/local/epilog` `-epiloglib=/usr/local/epilog/fe/lib`.

- `-epilogbin=<dir>`

Specifies the directory of where the Epilog binaries are to be found.

- `-epiloginc=<dir>`

Specifies the directory of where the epilog's included sources headers are to be found.

- `-MPITRACE`

Specifies the tracing option and generates event traces for MPI calls and routines that are ancestors of MPI calls in the callstack. This option is useful for generating traces that are converted to the EPI-LOG trace format. KOJAK's Expert automatic diagnosis tool needs traces with events that call MPI routines. Do not use this option with the -TRACE option.

- `-pythoninc=<dir>`

Specifies the location of the Python include directory. This is the directory where Python.h header file is located. This option enables python bindings to be generated. The user should set the environment variable PYTHONPATH to `<TAUROOT>/<ARCH>/lib/bindings-<options>` to use a specific version of the TAU Python bindings. By importing package pytau, a user can manually instrument the source code and use the TAU API. On the other hand, by importing tau and using `tau.run(<func>)`, TAU can automatically generate instrumentation. See examples/python directory for further information.

- `-pythonlib=<dir>`

Specifies the location of the Python lib directory. This is the directory where *.py and *.pyc files (and config directory) are located. This option is mandatory for IBM when Python bindings are used. For other systems, this option may not be specified (but `-pythoninc=<dir>` needs to be specified).

- `-PROFILESTATS`

Specifies the calculation of additional statistics, such as the standard deviation of the exclusive time/counts spent in each profiled block. This option is an extension of -PROFILE, the default profiling option.

- `-PROFILEMEMORY`

Specifies tracking heap memory utilization for each instrumented function. When any function entry takes place, a sample of the heap memory used is taken. This data is stored as user-defined event data in profiles/traces.

- `-PROFILEHEADROOM`

Specifies tracking memory available in the heap (as opposed to memory utilization tracking in -PROFILEMEMORY). When any function entry takes place, a sample of the memory available (headroom to grow) is taken. This data is stored as user-defined event data in profiles/traces. Please refer to the examples/headroom/README file for a full explanation of these headroom options and the C++/C/F90 API for evaluating the headroom.

- `-COMPENSATE`

Specifies online compensation of performance perturbation. When this option is used, TAU computes its overhead and subtracts it from the profiles. It can be only used when profiling is chosen. This option works with MULTIPLECOUNTERS as well, but while it is relevant for removing perturbation with wallclock time, it cannot accurately account for perturbation with hardware performance counts (e.g., L1 Data cache misses). See TAU Publication [Europar04] for further information on this option.

- `-PROFILECOUNTERS`

Specifies use of hardware performance counters for profiling under IRIX using the SGI R10000 prefix counter access interface. The use of this option is deprecated in favor of the `-pcl=<dir>` and `-papi=<dir>` options described above.

- `-muse`

Specifies the use of MAGNET/MUSE to extract low-level information from the kernel. To use this configuration, Linux kernel has to be patched with MAGNET and MUSE has to be installed on the executing machine. Also, magnetd has to be running with the appropriate handlers and filters installed. User can specify package by setting the environment variable TAU_MUSE_PACKAGE. MUSE [http://public.lanl.gov/radiant/]

- -noex

Specifies that no exceptions be used while compiling the library. This is relevant for C++.

- -useropt=<options-list>

Specifies additional user options such as -g or -I. For multiple options, the options list should be enclosed in a single quote. For example

```
% ./configure -useropt='-g -I/usr/local/stl'
```

- -help

Lists all the available configure options and quits.

1.1.5. tau_setup

tau_setup is a GUI interface to the configure and installtau tools.

1.1.6. installtau script

To install multiple (typical) configurations of TAU at a site, you may use the script `installtau'. It takes options similar to those described above. It invokes ./configure <opts>; make clean install; to create multiple libraries that may be requested by the users at a site. The installtau script accepts the following options:

```
% installtau -help

TAU Configuration Utility
*****
Usage: installtau [OPTIONS]
  where [OPTIONS] are:
  -arch=<arch>
  -fortran=<compiler>
  -cc=<compiler>
  -c++=<compiler>
  -useropt=<options>
  -pdt=<pdt_dir>
  -pdtcompdir=<compdir>
  -pdt_c++=<C++ Compiler>
  -papi=<papidir>
  -vtf=<vtfdir>
  -otf=<otfdir>
  -slog2=<dir> (for external slog2 dir)
  -slog2 (for using slog2 bundled with TAU)
  -dyninst=<dyninst_dir>
  -mpiinc=<mpiincdir>
  -mpilib=<mpilibdir>
```

```

-mpilibrary=<mpilibrary>
-perfinc=<dir>
-perflib=<dir>
-perflibrary=<library>
-mpi
-tag=<unique name>
-nocomm
-opari=<oparidir>
-epilog=<epilogdir>
-epiloginc=<absolute path to epilog include dir> (<epilog>/include default)
-epilogbin=<absolute path to epilog bin dir> (<epilog>/bin default)
-epiloglib=<absolute path to epilog lib dir> (<epilog>/lib default)
-prefix=<dir>
-exec-prefix=<dir>
*****

```

These options are similar to the options used by the configure script.

1.1.7. upgradetau

This script is provided to rebuild all TAU configurations previously built in a different TAU source directory. Give this command the location of a previous version of tau followed by any additional configurations and it will rebuild tau with these same options.

1.1.8. tau_validate

This script will attempt to validate a tau installation. it only argument is TAU's architecture directory. These are some options:

- -v Verbose output
- --html Output results in HTML
- --build Only build
- --run Only run

Here is a simple example:

```

bash : ./tau_validate --html x86_64 &> results.html
tcsh : ./tau_validate --html x86_64 >& results.html

```

1.2. Platforms Supported

TAU has been tested on the following platforms:

- 1. SGI

On IRIX 6.x based systems, including Indy, Power Challenge, Onyx, Onyx2 and Origin 200, 2000, 3000 Series, CC 7.2+, KAI [<http://www.kai.com>] KCC and g++ [<http://www.gnu.org>] compilers are supported. On SGI Altix systems, Intel, and GNU compilers are supported.

- 2. LINUX Clusters

On Linux based Intel x86 PC clusters, KAI/Intel's KCC, g++, egcs (GNU), pgCC (PGI) [<http://www.pgroup.com>], FCC (Fujitsu) [<http://www.fujitsu.com>] and icpc/ecpc Intel [<http://www.intel.com>] compilers have been tested. TAU also runs under IA-64, Opteron, PowerPC, Alpha, Apple PowerMac, Sparc and other processors running Linux.

- 3. Sun Solaris

Sun compilers (CC, F90), KAI KCC, KAP/Pro and GNU g++ work with TAU.

- 4. IBM

On IBM SP2 and AIX systems, KAI KCC, KAP/Pro, IBM x1C, xlc, xlf90 and g++ compilers work with TAU. On IBM BG/L, IBM x1C, blrts_x1C, blrts_xlf90, blrts_xlc, and gnu compilers work with TAU. On IBM pSeries Linux, x1C, xlc, xlf90 and gnu compilers work with TAU.

- 5. HP HP-UX

On HP PA-RISC systems, aCC and g++ can be used.

- 6. HP Alpha Tru64

On HP Alpha Tru64 machines, cxx and g++, and Guide compilers may be used with TAU.

- 7. NEC SX series vector machines

On NEC SX-5 systems, NEC c++ may be used with TAU.

- 8. Cray X1, T3E, SV-1, XT3, RedStorm, and Cray Compute Node Linux (CNL)

On Cray T3E systems, KAI KCC and Cray CC compilers have been tested with TAU. On Cray SV-1 and X1 systems, Cray CC compilers have been tested with TAU. On Cray XT3, and RedStorm systems, PGI and GNU compilers have been tested with TAU. For those using Cray CNL you need to configure tau with the option `-arch=craycnl`

- On Hitachi machines, Hitachi KCC, g++ and Hitachi cc compilers may be used with TAU

- 10. Apple OS X

On Apple OS X machines, c++ or g++ may be used to compile TAU. Also, IBM's xlf90, xlf and Absoft Fortran 90 compilers for G4/G5 may be used with TAU.

- 11. Fujitsu PRIMEPOWER

On Fujitsu Power machines, Sun and Fujitsu compilers may be used with TAU.

- 11. Microsoft Window

On Windows, Microsoft Visual C++ 6.0 or higher and JDK 1.2+ compilers have been tested with TAU

NOTE: TAU has been tested with JDK 1.2, 1.3, 1.4.x under Solaris, SGI, IBM, Linux, and MacOS X.

1.3. Software Requirements

- 1. Java v 1.4

TAU's GUI ParaProf requires Java v1.3 or better in your path. We recommend Java version 1.4x from Sun. An older Tcl/Tk based browser racy is also included with TAU for compatibility. It requires the executable wish to be in your path. Racy is also available in this distribution but support for racy will be gradually phased out. Users are encouraged to use ParaProf instead. ParaProf does *not* require -jdk=<dir> option to be specified (which is used for configuring TAU for analyzing Java applications). The Java program should be in the user's path.