

# Performance Tuning and Analysis Tools

## Objective

Become familiar with tools integrated with PTP, to help enhance performance of parallel applications

## Contents

Overview of ETFw and Performance Tools

# PTP/External Tools Framework

formerly "Performance Tools Framework"

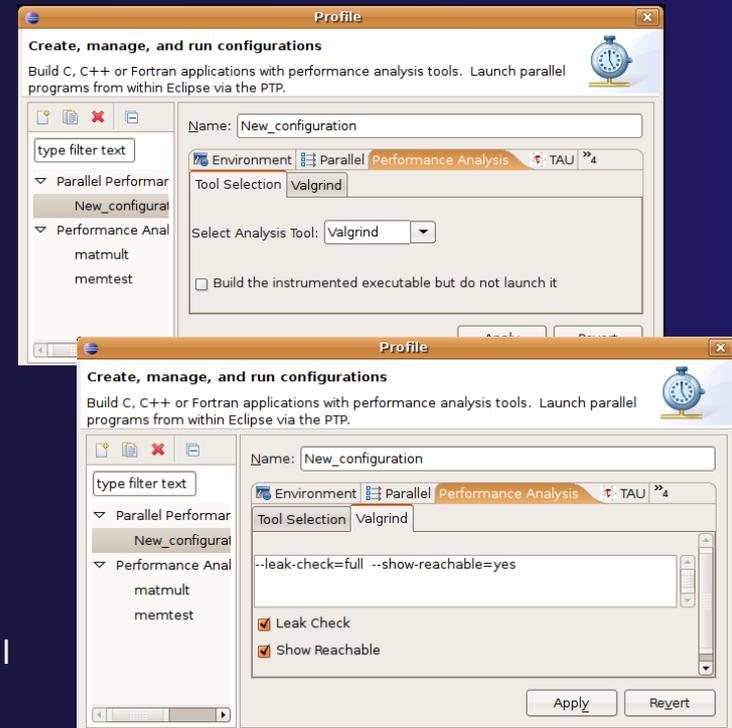
## Goal:

- ★ Reduce the "eclipse plumbing" necessary to integrate tools
- ★ Provide integration for instrumentation, measurement, and analysis for a variety of performance tools
  - ★ Dynamic Tool Definitions: Workflows & UI
  - ★ Tools and tool workflows are specified in an XML file
  - ★ Tools are selected and configured in the launch configuration window
  - ★ Output is generated, managed and analyzed as specified in the workflow
  - ★ One-click 'launch' functionality
  - ★ Support for development tools such as TAU, PPW and others.
  - ★ Adding new tools is much easier than developing a full Eclipse plug-in

```

-<tool name="Valgrind">
-<execute>
  <utility command="bash" group="inbin"/>
  -<utility command="valgrind" group="valgrind">
    -<optionpane title="Valgrind" separatewith=" ">
      <togoption label="Leak Check" optname="--leak-check=full" tooltip="Leak Check" />
      <togoption label="Show Reachable" optname="--show-reachable=yes" tooltip="Show Reachable" />
    </optionpane>
  </utility>
</execute>
</tool>

```



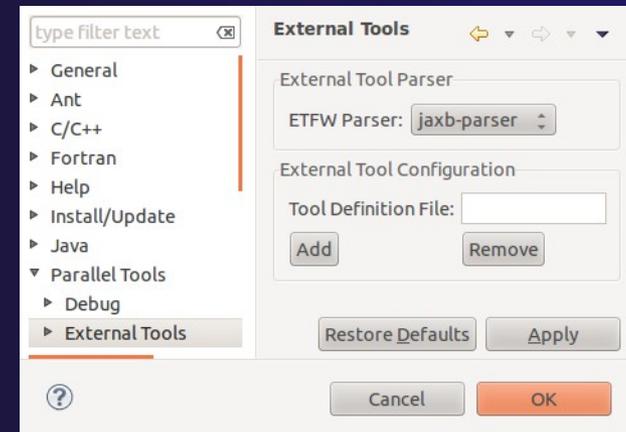
# SAX and JAXB Tool Definitions

Prior implementations of ETFW used a simple SAX based schema to define tool workflows

By default workflows now use the more powerful JAXB schema that defines PTP's resource manager

Legacy workflows can still be loaded by selecting the SAX parser in PTP options

Window->Preferences->  
Parallel Tools->External Tools



# Performance Tuning and Analysis Tools - TAU

## ★ Objective

- ★ Become familiar with tools integrated with PTP, to help enhance performance of parallel applications

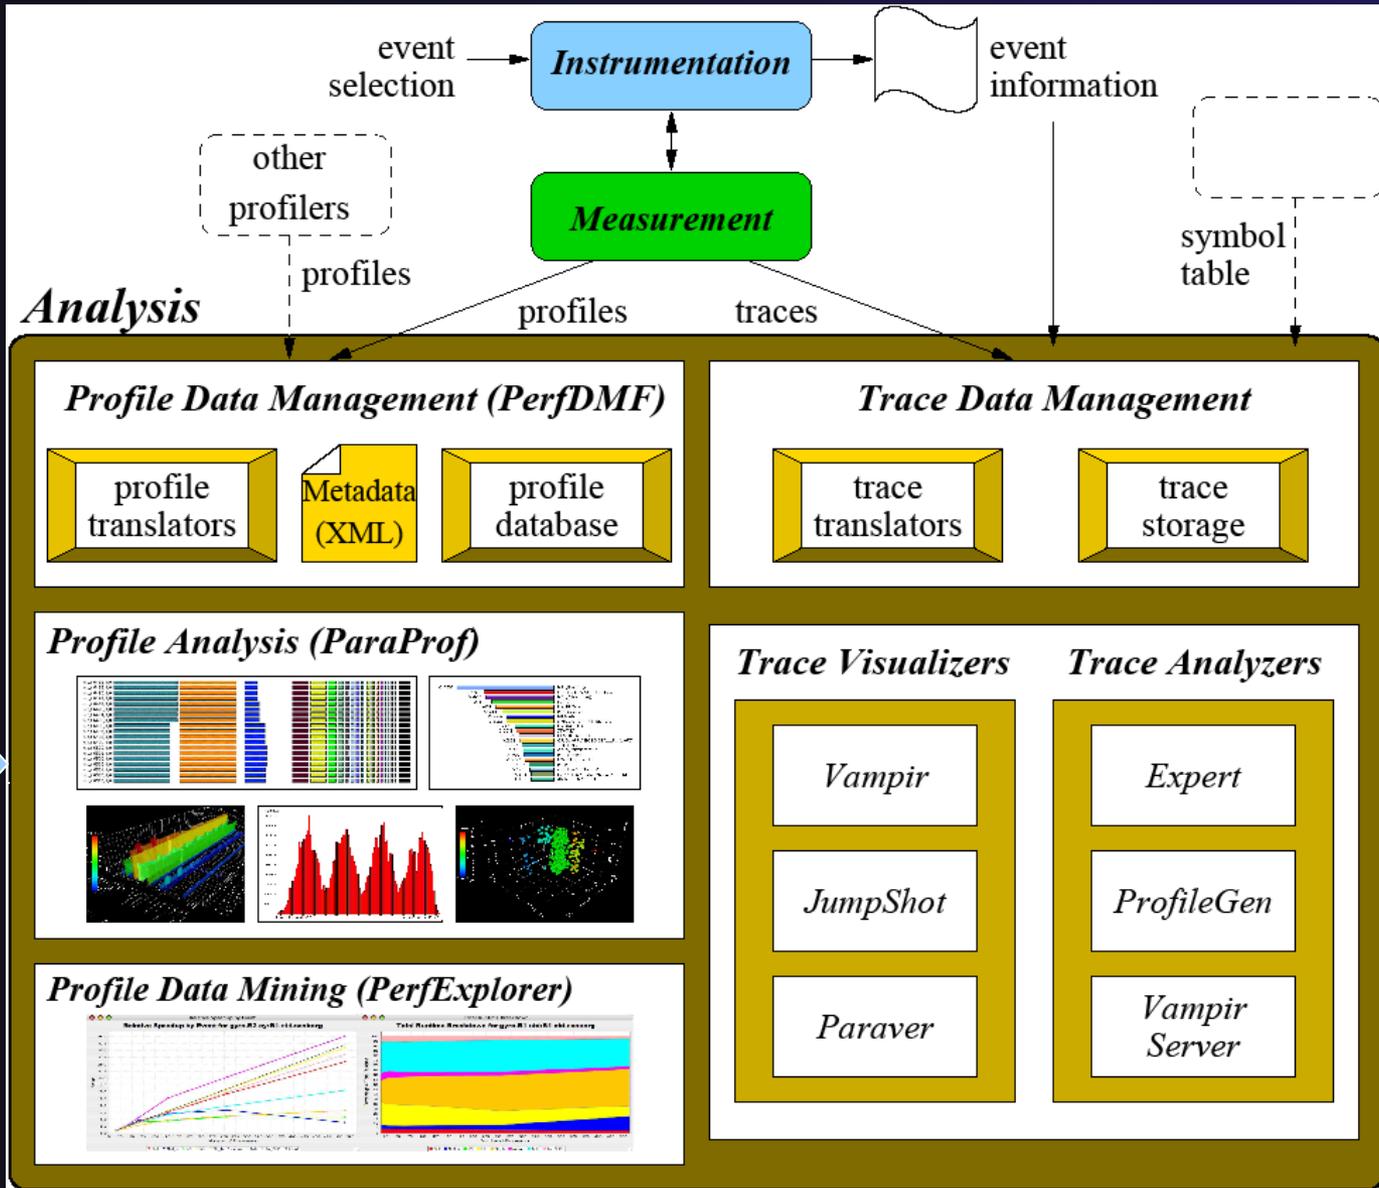
## ★ Contents

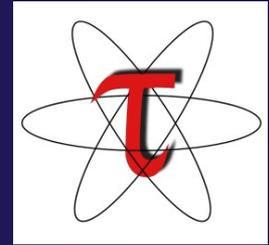
- ★ Performance Tuning and external tools:
  - ★ PTP External Tools Framework (ETFw), TAU
  - Hands-on exercise using TAU with PTP

# TAU: Tuning and Analysis Utilities

- ★ TAU is a performance evaluation tool
- ★ It supports parallel profiling and tracing
  - ★ Profiling shows you how much (total) time was spent in each routine
  - ★ Tracing shows you *when* the events take place in each process along a timeline
- ★ TAU uses a package called PDT (Performance Database Toolkit) for automatic instrumentation of the source code
- ★ Profiling and tracing can measure time as well as hardware performance counters from your CPU (or GPU!)
- ★ TAU can automatically instrument your source code (routines, loops, I/O, memory, phases, etc.)
- ★ TAU runs on all HPC platforms and it is free (BSD style license)
- ★ TAU has instrumentation, measurement and analysis tools
  - ★ **paraprof** is TAU's 3D profile browser

# TAU Performance System Architecture





# PTP TAU plug-ins

<http://www.cs.uoregon.edu/research/tau>

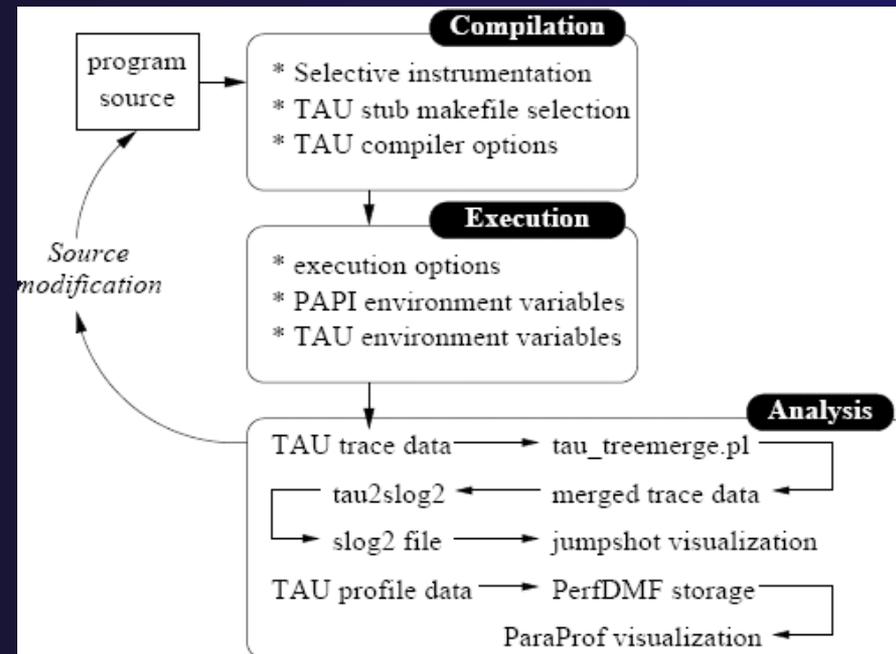
- ★ TAU (Tuning and Analysis Utilities)
- ★ First implementation of External Tools Framework (ETFw)
- ★ Eclipse plug-ins wrap TAU functions, make them available from Eclipse
- ★ Full GUI support for the TAU command line interface
- ★ Performance analysis integrated with development environment

The screenshots illustrate the TAU workflow in Eclipse:

- Configuration:** The 'Create, manage, and run configurations' dialog allows users to create a configuration for launching a program to be instrumented and profiled by TAU.
- Counter Selection:** The 'PAPI Counters' dialog allows users to select the PAPI counters to use with TAU.
- MPI Configuration:** The 'MPI' configuration options allow users to select the MPI implementation and the PAPI counters to use.
- Analysis:** The 'MPI Barrier' code is analyzed, and the results are visualized in a 3D plot and a list of performance metrics.

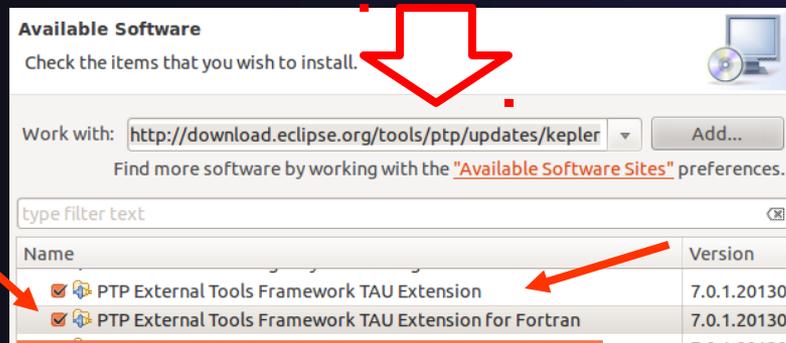
# TAU Integration with PTP

- ★ TAU: Tuning and Analysis Utilities
  - ★ Performance data collection and analysis for HPC codes
  - ★ Numerous features
  - ★ Command line interface
- ★ The TAU Workflow:
  - ★ Instrumentation
  - ★ Execution
  - ★ Analysis



# TAU PTP Installation

- ✦ This tutorial assumes that the TAU extensions for PTP are installed – they are not included in the “Eclipse for Parallel Application Developers”
- ✦ The installation section (Module 1) shows how to install TAU and other features from the PTP update site – be sure TAU was selected



To confirm:

- ✦ Help>Install New Software...
- ✦ Select the link “What is already installed” at the bottom of the dialog
- ✦ You should see the TAU Extension

# Installing TAU Analysis Tools

- ★ The TAU plugin can use ParaProf for visual analysis and TauDB for organization of profiles
- ★ To install these utilities on Mac or Linux platforms:
  - ★ Download (browser, curl or wget)  
[tau.uoregon.edu/tautools.tgz](http://tau.uoregon.edu/tautools.tgz)
  - ★ `tar -zxf tautools.tgz`
  - ★ `cd tautools-2.22.2`
  - ★ `./configure`
  - ★ Set path as shown (launch eclipse from this environment)
  - ★ Run `taudb_configure` and follow the instructions
- ★ Java WebStart: [tau.uoregon.edu/paraprof](http://tau.uoregon.edu/paraprof)
- ★ TAU Installation, downloads and instructions: [tau.uoregon.edu](http://tau.uoregon.edu)

# Assumptions

- ✦ Obtain and install TAU\*
  - ✦ Download at [tau.uoregon.edu](http://tau.uoregon.edu)
  - ✦ The website includes setup and user guides
- ✦ Set up the `$PATH` on the remote machine\*
  - ✦ For TAU you should be able to run `'which pprof'` on a remote login and see a result from your TAU bin directory
  - ✦ On [trestles.sdsc.edu](http://trestles.sdsc.edu) this is accomplished by loading the tau module in the environment manager for the build and launch configurations
- ✦ Include `'eclipse.inc'` in the makefile\*
  - ✦ Create an empty `eclipse.inc` file in the same directory as the makefile
  - ✦ Place `'include eclipse.inc'` in the makefile after regular compiler definitions
  - ✦ ETFw will modify `eclipse.inc` to set `CC/CXX/FC` variables

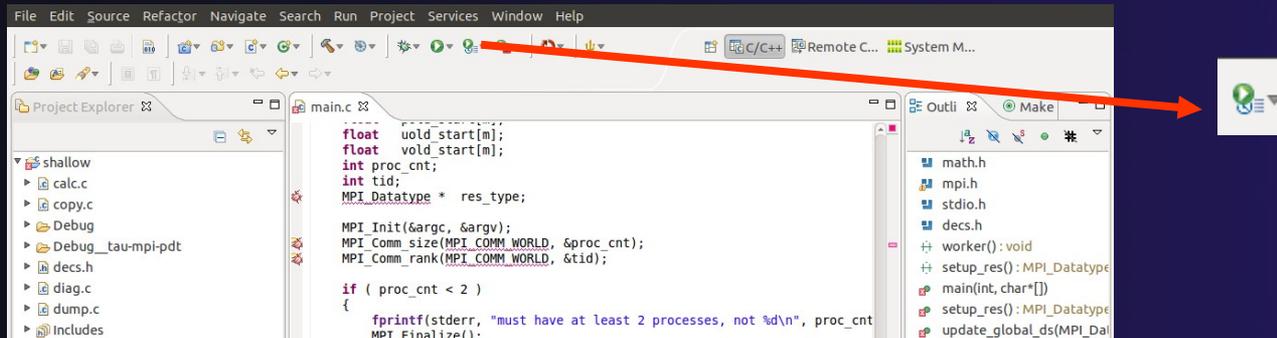
# Selective Instrumentation

- ✦ By default tau provides timing data for each subroutine of your application
- ✦ Selective instrumentation allows you to include/exclude code from analysis and control additional analysis features
  - ✦ Include/exclude source files or routines
  - ✦ Add timers and phases around routines or arbitrary code
  - ✦ Instrument loops
  - ✦ Note that some instrumentation features require the PDT
- ✦ Right click on `calc.c`, `init.c`, `diag.c` go to the Selective Instrumentation option and select Instrument Loops
- ✦ Note the creation of `tau.selective` (refresh if needed)

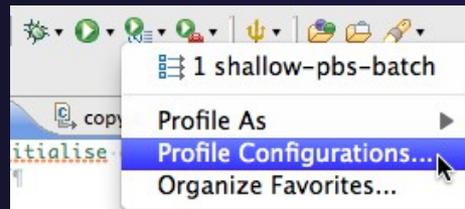


# Begin Profile Configuration

- ✦ The ETFw uses the same run configurations and resource managers as debugging/launching
- ✦ Click on the 'Run' menu or the right side of the Profile button



- ✦ From the dropdown menu select 'Profile configurations...'





# Select Configuration

- ✦ Select the shallow configuration prepared earlier
- ✦ The Resource and Application configuration tabs require little or no modification
  - ✦ We are using the same resource manager and Torque settings
  - ✦ Since we are using a makefile project the application will be rebuilt in and run from the previously selected location

Performance Analysis tab is present in the **Profile Configurations** dialog

Create, manage, and run configurations  
[Performance Analysis]: No workflow selected.

Name: shallow-xse13

Resources Application Arguments Environment Performance Analysis Common Parametric Study

Target System Configuration: edusdsc.trestles.torque.batch

Connection Type  
 Local  Remote trestles

Basic Settings Advanced Settings Import Script

| Name                     | Value                                       | Description  |
|--------------------------|---|--|
| Job Name:                | ptp_job                                     | The name assigned to the job by the qsub or qalter command.  |
| Account:                 |   | Account to which to charge this job.   |
| Queue:                   | shared                                      | Designation of the queue to which to submit the job.   |
| Number of nodes:         | 1:ppn=5                                     | Number and/or type of nodes to be reserved for exclusive use by the job. [usage hint] number_nodes:ppn=N |
| Total Memory Needed:     |   | Maximum amount of memory used by all concurrent processes in the job.                                    |
| Wallclock Time:          | 00:30:00                                    | Maximum amount of real time during which the job can be in the running state.                            |
| MPI Command:             | mpirun                                      | Which mpi command to use.  |
| MPI Number of Processes: | 5   | the '-np' value [usually equals Nodes*ppn]   |
| Export Environment:      | <input checked="" type="checkbox"/>         | All variables in the qsub command's environment are to be exported to the b.                             |
| Modules to Load:         | <input type="button" value="Configure..."/> | Modules that will be loaded inside the job script.   |

View Script View Configuration Restore Defaults

Filter matched 4 of 4 items

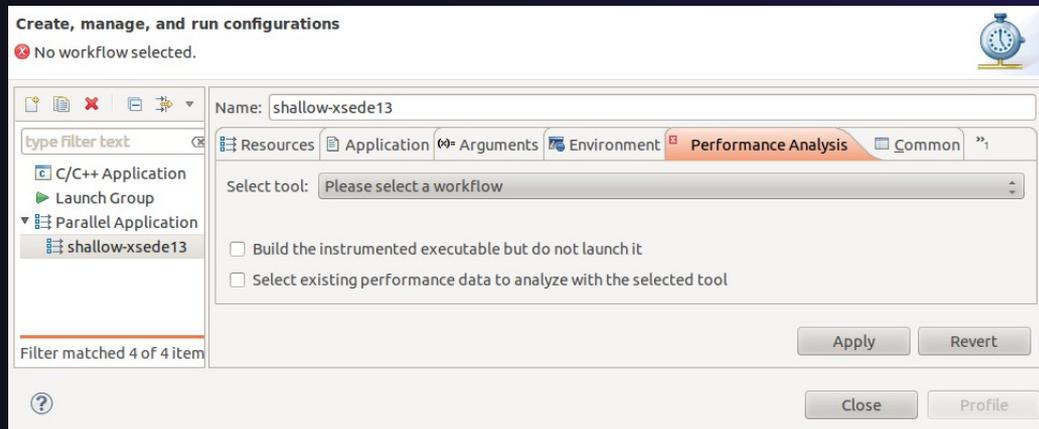
Apply Revert

Close Profile



# Select Tool/Workflow

- ✦ Select the **Performance Analysis** tab and choose the TAU tool set in the 'Select Tool' dropdown box
  - ✦ Other tools may be available, either installed as plug-ins or loaded from workflow definition XML files
  - ✦ Configuration sub-panes appear depending on the selected tool



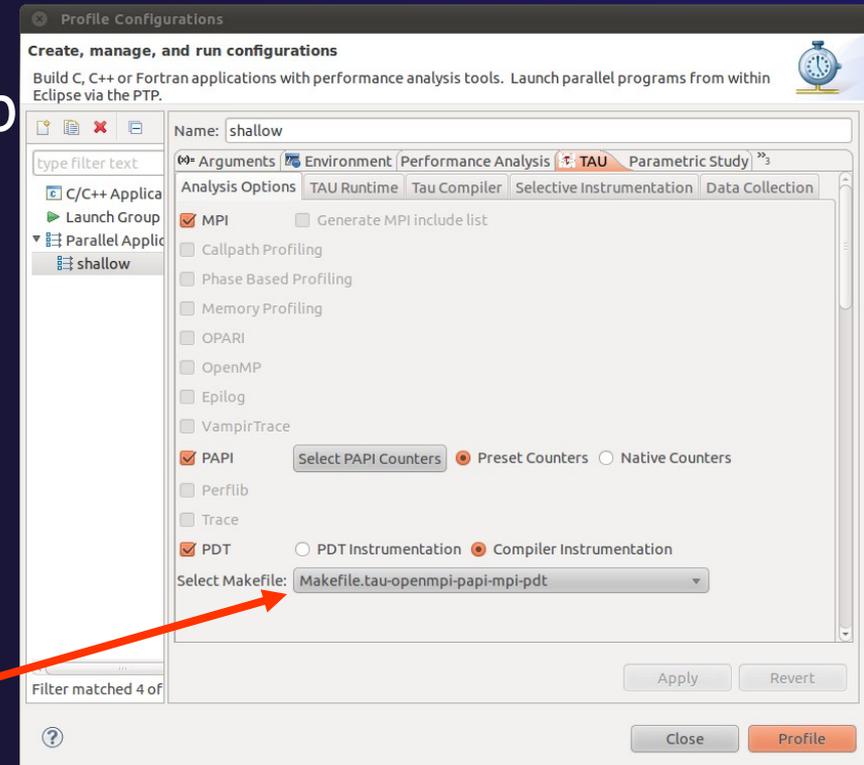
Tabs may be hidden if the window is too small



# Select TAU Configuration

## ★ Choose the TAU Makefile tab

- ★ All TAU configurations in remote installation are available
- ★ Check MPI and PDT checkboxes to filter listed makefiles
- ★ Make your selection in the **Select Makefile:** dropdown box
- ★ Select Makefile.tau-mpi-pdt





# Choose PAPI Hardware Counters

★ When a PAPI-enabled TAU configuration is selected the PAPI Counter tool becomes available



★ Select the 'Select PAPI Counters' button to open the tool

★ Open the PRESET subtree

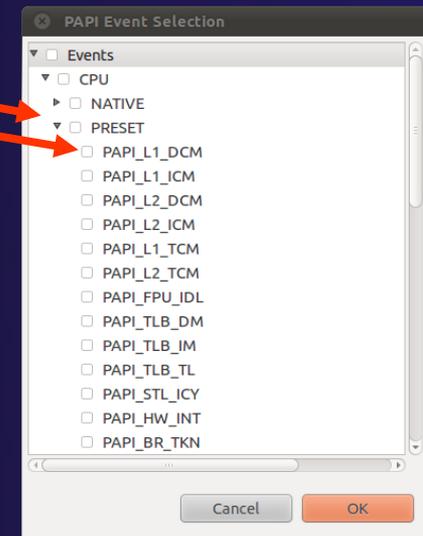
★ Select PAPI\_L1\_DCM (Data cache misses)

★ Scroll down to select PAPI\_FP\_INS (Floating point instructions)

★ Invalid selections are automatically excluded

★ Select **OK**

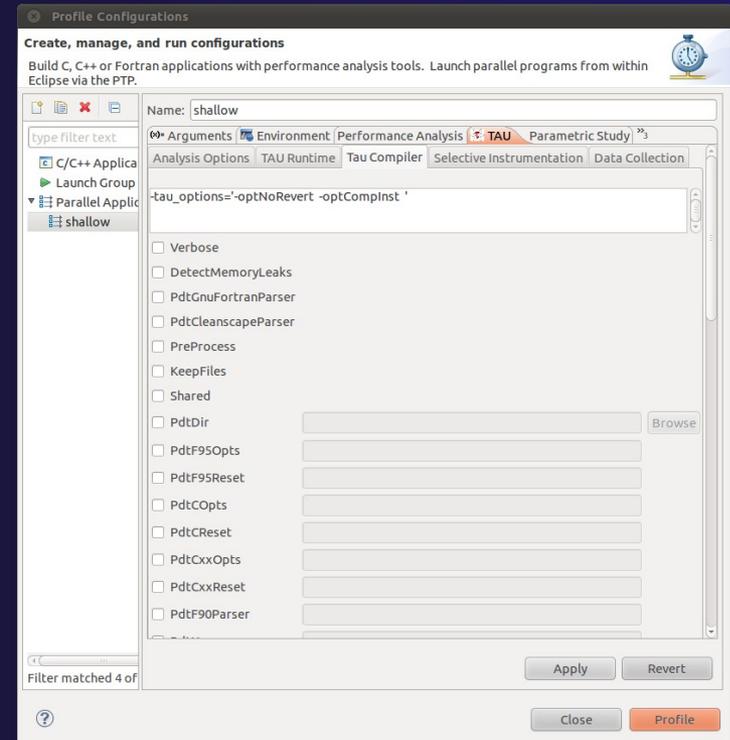
★ **Not available on trestles.sdsc.edu**





# Compiler Options

- ✦ TAU Compiler Options
  - ✦ Set arguments to TAU compiler scripts
  - ✦ Control instrumentation and compilation behavior
  - ✦ Verbose shows activity of compiler wrapper
  - ✦ KeepFiles retains instrumented source
  - ✦ PreProcess handles C type ifdefs in fortran
- ✦ In the Selective Instrumentation tab select Internal then hit Apply
- ✦ Scroll to bottom of the Tau Compiler tab and activate TauSelectFile to use tau.selective

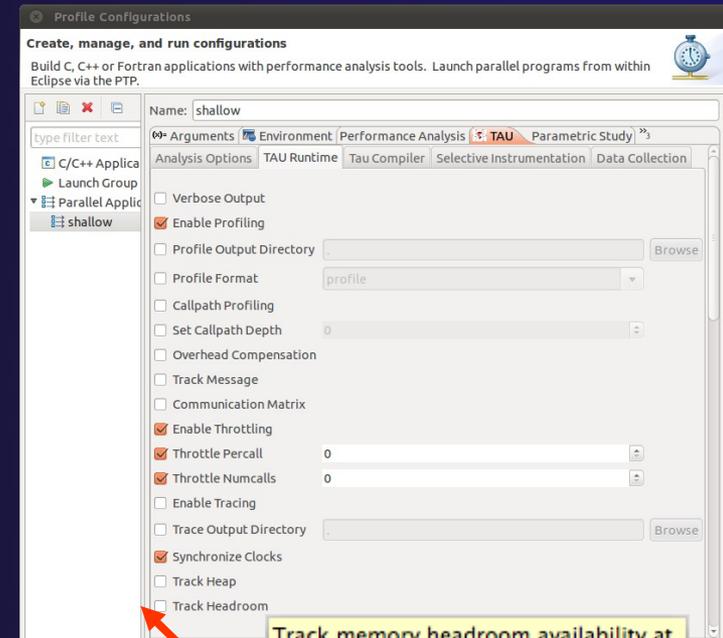




# Runtime Options

## ✦ TAU Runtime options

- ✦ Set environment variables used by TAU
- ✦ Control data collection behavior
- ✦ Verbose provides debugging info
- ✦ Callpath shows call stack placement of events
- ✦ Throttling reduces overhead
- ✦ Tracing generates execution timelines
- ✦ Set Profile Format to merged



Hover help



# Working with Profiles

- ✦ Profiles are uploaded to selected database
- ✦ A text summary may be printed to the console
- ✦ Profiles may be uploaded to the TAU Portal for viewing online
  - ✦ tau.nic.uoregon.edu
- ✦ Profiles may be copied to your workspace and loaded in ParaProf from the command line. Select Keep Profiles

Analysis Options | TAU Runtime | Tau Compiler | Selective Instrumentation | Data Collection

Select Database:

Keep profiles

Print Profile Summary

Upload profile data to TAU Portal

Console | Properties | Problems | Tasks

TAU Profile Output

MULTI\_GET\_TIME\_OF\_DAYReading Profile files in profile.\*

FUNCTION SUMMARY (total):

| %Time | Exclusive msec | Inclusive total msec | #Call | #Subrs | Inclusive Name usec/call |
|-------|----------------|----------------------|-------|--------|--------------------------|
| 100.0 | 196            | 16,797               | 9     | 9      | 1866428 .TAU application |
| 98.8  | 56             | 16,601               | 9     | 3662   | 1844640 main             |
| 68.7  | 11,538         | 11,538               | 9     | 0      | 1282002 MPI_Init()       |
| 26.0  | 204            | 4,360                | 8     | 59684  | 545009 worker            |
| 9.5   | 803            | 1,602                | 80000 | 160000 | 20 neighbour_receive     |
| 8.3   | 789            | 1,402                | 80000 | 160000 | 18 neighbour_send        |
| 8.3   | 234            | 1,398                | 8000  | 64000  | 175 time_load            |
| 7.1   | 1,188          | 1,188                | 81968 | 0      | 14 MPI_Recv()            |
| 6.5   | 182            | 1,085                | 8000  | 48000  | 136 calc_load            |

TAU Portal URL:

TAU Portal Username:

TAU Portal Password:

Select a workspace

Select Workspace:

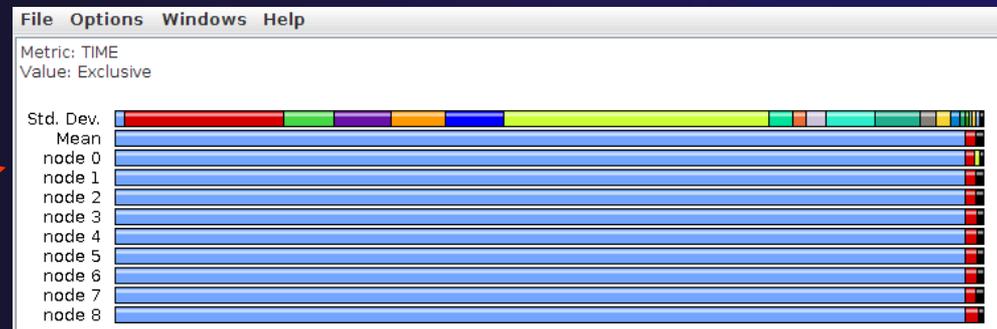


# Launch TAU Analysis

- ★ Once your TAU launch is configured select 'Profile'
- ★ Notice that the project rebuilds with TAU compiler commands
- ★ The project will execute normally but TAU profiles will be generated
- ★ TAU profiles will be processed as specified in the launch configuration.
- ★ If you have a local profile database the run will show up in the Performance Data Management view
  - ★ Double click the new entry to view in ParaProf
  - ★ Right click on a function bar and select **Show Source Code** for source callback to Eclipse



TAU

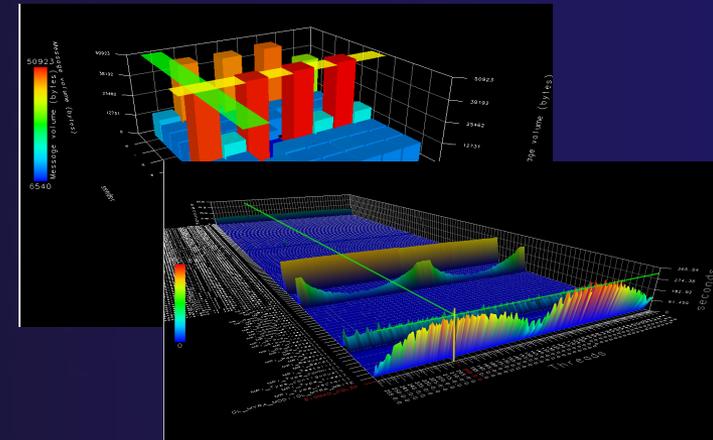


TAU-21



# Paraprof

- ✦ Use ParaProf for profile visualization to identify performance hotspots
  - ✦ Inefficient sequential computation
  - ✦ Communication overhead
  - ✦ IO/Memory bottlenecks
  - ✦ Load imbalance
  - ✦ Suboptimal cache performance
- ✦ Compare multiple trials in PerfExplorer to identify performance regressions and scaling issues
- ✦ To use ParaProf, install TAU from [tau.uoregon.edu](http://tau.uoregon.edu) or use Java webstart from [tau.uoregon.edu/paraprof](http://tau.uoregon.edu/paraprof)





# Exercise

- ★ Multi-Trial profile comparison
  1. Edit the shallow Makefile, adding -O3 to CFLAGS and FFLAGS
  2. Rerun the analysis (Run->Profile Configurations. Hit Profile)
  3. A second trial, distinguished by a new timestamp, will be generated
    - ★ It will appear in your Performance Data Manager view if a profile database is available
    - ★ Also present in the Profile subdirectory of your project directory
    - ★ If you do not see a Profile directory right click on your project and go to Synchronization->'Sync All Now'
  4. Load the two trials in paraprof (on the command line: paraprof /path/to/tauprofile.xml)
  5. Open Windows->ParaProf Manager
  6. Expand your database down to reveal all trials
  7. Right click on each trial and click 'Add Mean to Comparison Window' to visualize the two trials side by side