

PROFILING CFS CODE

```
roshan@login01:>module load perftools
roshan@login01:>
```

- Makes the default version of CrayPAT available
- Subsequent compiler invocations will automatically insert necessary hooks for profiling (not always up-to-date with latest third-party compilers)
- Binaries are *not* automatically instrumented

```
roshan@login01:>module list
Currently Loaded Modulefiles:
 1) modules/3.2.10.6
 2) eproxy/2.0.16-6.0.4.1_3.1__g001b199.ari
 3) cce/8.6.3
 4) craype-network-aries
 5) craype/2.5.13
 6) cray-libsci/17.09.1
 7) udreg/2.3.2-6.0.4.0_12.2__g2f9c3ee.ari
 8) ugni/6.0.14-6.0.4.0_14.1__ge7db4a2.ari
 9) pmi/5.0.12
10) dmapp/7.1.1-6.0.4.0_46.2__gb8abda2.ari
11) gni-headers/5.0.11-6.0.4.0_7.2__g7136988.ari
12) xpmem/2.2.2-6.0.4.1_18.1__g43b0535.ari
13) job/2.2.2-6.0.4.0_8.2__g3c644b5.ari
14) dvs/2.7_2.2.36-6.0.4.1_16.3__g4c8274a
15) alps/6.4.3-6.0.4.1_2.1__g92a2fc0.ari
16) rca/2.2.11-6.0.4.0_13.2__g84de67a.ari
17) atp/2.1.1
18) perftools-base/6.5.2
19) PrgEnv-cray/6.0.4
20) cray-mpich/7.6.3
21) pbs/default
22) perftools
roshan@login01:>
```

- Build the binary of the application/program/code with “pat_build” command you want to profile. Here, we are trying to profile cfs_mlc_coupler binary.

```
roshan@login01:>
roshan@login01:>pat_build cfs_mlc_coupler
```

NOTE : If you are getting below error while executing the above command :

```

roshan@elogin01:>pat_build cfs_mlc_coupler
ERROR: Missing required ELF section '.note.link' from the program '/home/rtayde/cFs/cfs.v2.1.18-rfp/exec/cfs_mlc_coupler'. Load the
correct 'perftools' module and rebuild the program.
roshan@elogin01:>
roshan@elogin01:>
roshan@elogin01:>

```

Resolution : You have to rebuild the binary with the perftools module loaded.

Ex. :

```

roshan@elogin01:>ls
cfs_mlc_coupler  coupler.o      COUPLER_MODULE.mod  Makefile  coupler.f      makefile.sh.orig
makefile.sh     coupler_module.o  cfs_mlc_coupler_working  build-ao  coupler_module.f
roshan@elogin01:>pat_build cfs_mlc_coupler
ERROR: Missing required ELF section '.note.link' from the program '/home/rtayde/CFS/cray/cfs.v2.1.18-rfp/src/cfs_ao_coupler.fd/cfs
_mlc_coupler'. Load the correct 'perftools' module and rebuild the program.
roshan@elogin01:>rm cfs_mlc_coupler
roshan@elogin01:>sh makefile.sh 2> /dev/null 1> /dev/null
roshan@elogin01:>pat_build cfs_mlc_coupler
roshan@elogin01:>ls
cfs_mlc_coupler+pat  makefile.sh  coupler_module.o  cfs_mlc_coupler_working  build-ao  coupler_module.f
cfs_mlc_coupler     coupler.o   COUPLER_MODULE.mod  Makefile                 coupler.f  makefile.sh.orig
roshan@elogin01:>

```

If the command is successful <binary name>+pat file will get created. In our example cfs_mlc_coupler+pat is got created.

Now lets try to profile cfs_atmos :

```

roshan@elogin01:>pat_build -S cfs_atmos_fcst
roshan@elogin01:>

```

- **-S** option Builds code with profiling hooks, then instruments the binary.
- Result named cfs_atmos_fcst+pat
- To know more about available options, please execute : man pat_build command.
- After executing above command below directory will be created in the output directory of the application :

NOTE :

- Running the “+pat” binary creates a data file “*.xf” or a directory in run directory
- pat_report reads that data file and prints lots of human-readable performance data. Creates an *.ap2 file.

```

roshan@eLogin01:>ls -ld cfs_atmos_fcst+48966-1503s/
drwxr-x--- 3 rlayde appsupp 4096 Aug  2 12:23 cfs_atmos_fcst+48966-1503s/
roshan@eLogin01:>
roshan@eLogin01:>

```

NOTE : THE NAME OF THE OUTPUT DIRECTORY MAY VARY.

To generate profile report execute below command :

pat_report <profiling directory name>

```

roshan@eLogin01:>pat_report cfs_atmos_fcst+48966-1503s
Skipping file because of suffix: cfs_atmos_fcst+48966-1503s/xf-files/001660.ap2
xf_ap2: No input binary data files were found.
pat_report: Error in creating cfs_atmos_fcst+48966-1503s/ap2-files/001660.ap2.ap2.
roshan@eLogin01:>pat_report cfs_atmos_fcst+48966-1503s*
Skipping file because of suffix: cfs_atmos_fcst+48966-1503s/xf-files/001660.ap2
xf_ap2: No input binary data files were found.
pat_report: Error in creating cfs_atmos_fcst+48966-1503s/ap2-files/001660.ap2.ap2.
roshan@eLogin01:>ls cfs_atmos_fcst+48966-1503s/
xf-files
roshan@eLogin01:>

```

As you can see its unable to find the *.ap2 files, we need to generate it. [NOTE : you may get such errors sometime.]

Resolution : you need to generate ap2 file.

```

roshan@eLogin01:>cd cfs_atmos_fcst+48966-1503s/xf-files/
roshan@eLogin01:>
roshan@eLogin01:>
roshan@eLogin01:>pat_report -f ap2 *.xf
Output redirected to: 001660.ap2
roshan@eLogin01:>
roshan@eLogin01:>
roshan@eLogin01:>
roshan@eLogin01:>
roshan@eLogin01:>ls
001660.ap2      001660.xf  001662.xf  001664.xf  001666.xf  001668.xf
build-options.apa 001661.xf  001663.xf  001665.xf  001667.xf  001669.xf
roshan@eLogin01:>

```

As you can see 001660.ap2 file is now created and which you can use to analyze the profile data.

```
roshan@e login01:>  
roshan@e login01:>pat_report -O profile+src 001660.ap2 █
```

The above command is used to show the profile data and the source code.

To know more please execute :

man pat_report

Or

pat_report -O help

Sample Output :

```

100.0% | 16,644.4 | -- | -- | Total
-----
39.2% | 6,517.3 | -- | -- | USER
-----
3 | 8.3% | 1,387.8 | -- | -- | swflux$spcvrt$module_radsw_main_
4 | 1.0% | 172.3 | 78.7 | 31.7% | cfs.v2.1.18-rfp/sorc/cfs_atmos_fcst.fd/radsw_main.f
4 | 4.0% | 668.0 | -- | -- | line.2743
3 | 2.8% | 463.0 | -- | -- | rtrnmc$module_radlw_main_
3 | 2.3% | 380.8 | -- | -- | cfs.v2.1.18-rfp/sorc/cfs_atmos_fcst.fd/radlw_main.f
3 | 1.8% | 292.8 | -- | -- | mcica_subcol$module_radlw_main_
3 | 1.2% | 204.9 | -- | -- | cfs.v2.1.18-rfp/sorc/cfs_atmos_fcst.fd/radlw_main.f
3 | 1.1% | 178.4 | -- | -- | random_number_t$mersenne_twister_
3 | 1.1% | 178.4 | -- | -- | cfs.v2.1.18-rfp/sorc/cfs_atmos_fcst.fd/mersenne_twister.f
3 | 1.1% | 178.4 | -- | -- | radb4_
3 | 1.1% | 178.4 | -- | -- | cfs.v2.1.18-rfp/sorc/cfs_atmos_fcst.fd/fftpack.f
3 | 1.1% | 178.4 | -- | -- | dcrfft_
3 | 1.1% | 178.4 | -- | -- | cfs.v2.1.18-rfp/sorc/cfs_atmos_fcst.fd/fftpack.f
3 | 1.1% | 178.4 | -- | -- | mcica_subcol$module_radsw_main_
3 | 1.1% | 178.4 | -- | -- | cfs.v2.1.18-rfp/sorc/cfs_atmos_fcst.fd/radsw_main.f
=====
29.7% | 4,939.6 | -- | -- | MPI
-----
12.3% | 2,045.5 | 885.5 | 30.6% | MPID_nem_gni_poll
9.0% | 1,489.9 | 10,899.1 | 89.0% | MPIDI_Cray_shared_mem_coll_bcst
4.3% | 720.5 | 313.5 | 30.7% | MPIDI_CH3I_Progress
2.5% | 423.9 | 181.1 | 30.3% | MPID_nem_gni_check_localCQ
=====
20.0% | 3,329.1 | -- | -- | GNI
-----

```

Another Example :

In the below example we are trying to trace MPI routines.

```

roshan@login01:~$ pat_build -u -g mpi cfs_atmos_fcst

```

Sample Output After Successful Execution:

```

roshan@ellogin01:>
roshan@ellogin01:>
roshan@ellogin01:>pat_report -O caller+src cfs_atmos_fcst+57414-1503t/
CrayPat/X: Version 6.5.2 Revision ba33e9b 08/22/17 16:05:58

```

```

Number of PEs (MPI ranks): 163 total, 90 with data in .ap2
Numbers of PEs per Node: 9 PEs on each of 10 Nodes with data
Numbers of Threads per PE: 1 thread on 1 PE
                          3 threads on each of 89 PEs
Number of Cores per Socket: 18
Execution start time: Tue Jul 31 15:11:25 2018
System name and speed: nid01670 2101 MHz (approx)
Intel Broadwell CPU Family: 6 Model: 79 Stepping: 1

```

```

-----
| 39.1% | 83.657033 | 57,573.2 | MPI
-----
| 32.3% | 69.075686 | 3,121.7 | MPI_RECV
3 | 32.3% | 69.028011 | 2,816.0 | disassemble_cc_:tiles.f:line.195
4 | | | | | atm_tiles_recv$atm_getsstice_:atm.comm.f:line.838
5 | 31.4% | 67.043975 | 704.0 | atm_getsstice_:atm.comm.f:line.738
6 | | | | | do_tstep_:dotstep_sicdif.f:line.269
-----
7 | ||||| 24.4% | 52.065164 | 686.4 | gfs_run$gfs_run_esmfmod_:GFS_Run_ESMFMod.f:line.25
8 | ||||| | | | | | run$gfs_gridcomp_esmfmod_:GFS_GridComp_ESMFMod.f:line.579
9 | ||||| | | | | | ESMC_FTable::ESMC_FTableCallVFuncPtr:ESMC_FTable.C:line.854
10 | ||||| | | | | | esmf_compexecute$esmf_compmo_:ESMF_Comp.F90:line.901
11 | ||||| | | | | | esmf_gridcomprun$esmf_gridcompmo_:ESMF_GridComp.F90:line.881
12 | ||||| | | | | | gfs_run_:Ensemble_sub.f:line.131
13 | ||||| | | | | | gfs_standalone_esmf_:GFS_Standalone_ESMF_ENS.f:line.319
7 | ||||| 7.0% | 14.978812 | 17.6 | tldfi_:digifilt.f:line.1
8 | ||||| | | | | | gfs_run$gfs_run_esmfmod_:GFS_Run_ESMFMod.f:line.25
9 | ||||| | | | | | run$gfs_gridcomp_esmfmod_:GFS_GridComp_ESMFMod.f:line.579
10 | ||||| | | | | | ESMC_FTable::ESMC_FTableCallVFuncPtr:ESMC_FTable.C:line.854
11 | ||||| | | | | | esmf_compexecute$esmf_compmo_:ESMF_Comp.F90:line.901
12 | ||||| | | | | | esmf_gridcomprun$esmf_gridcompmo_:ESMF_GridComp.F90:line.881
13 | ||||| | | | | | gfs_run_:Ensemble_sub.f:line.131
14 | ||||| | | | | | gfs_standalone_esmf_:GFS_Standalone_ESMF_ENS.f:line.319
-----
||| 4.1% | 8.801043 | 3,770.6 | mpi_alltoallv
-----
3 | 1.6% | 3.337664 | 1,459.6 | four2fln_:four2fln.f:line.187
3 | 1.1% | 2.291718 | 851.4 | sumfln_r_:sumfln.f:line.1010
-----
|| 2.1% | 4.384059 | 16,096.2 | mpi_send
3 | 1.2% | 2.510114 | 12,674.9 | assemble_cc_:tiles.f:line.115
4 | 1.2% | 2.502273 | 12,672.0 | atm_sendflux$atm_sendfluxes_:atm.comm.f:line.1197

```

For further help/queries please mail us at :

PRATYUSHSUPPORT@TROPMET.RES.IN