


VI-HPS



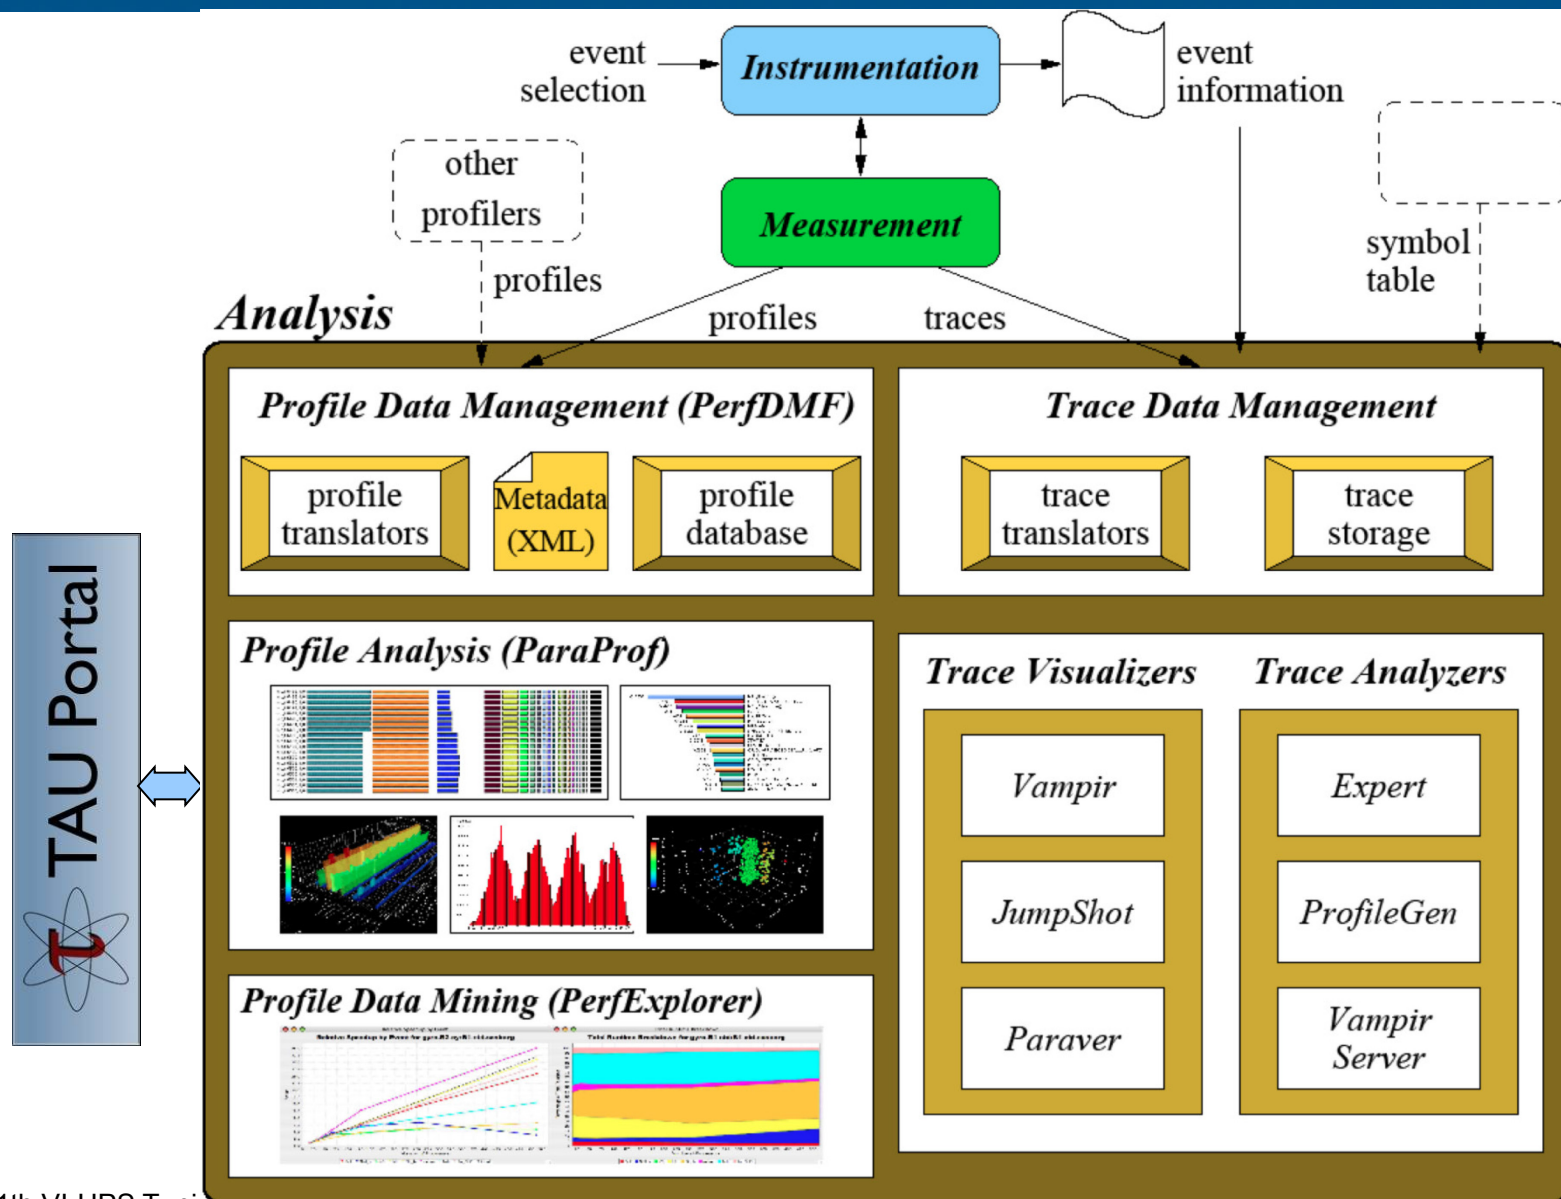
Profile Data Mining with PerfExplorer



Sameer Shende

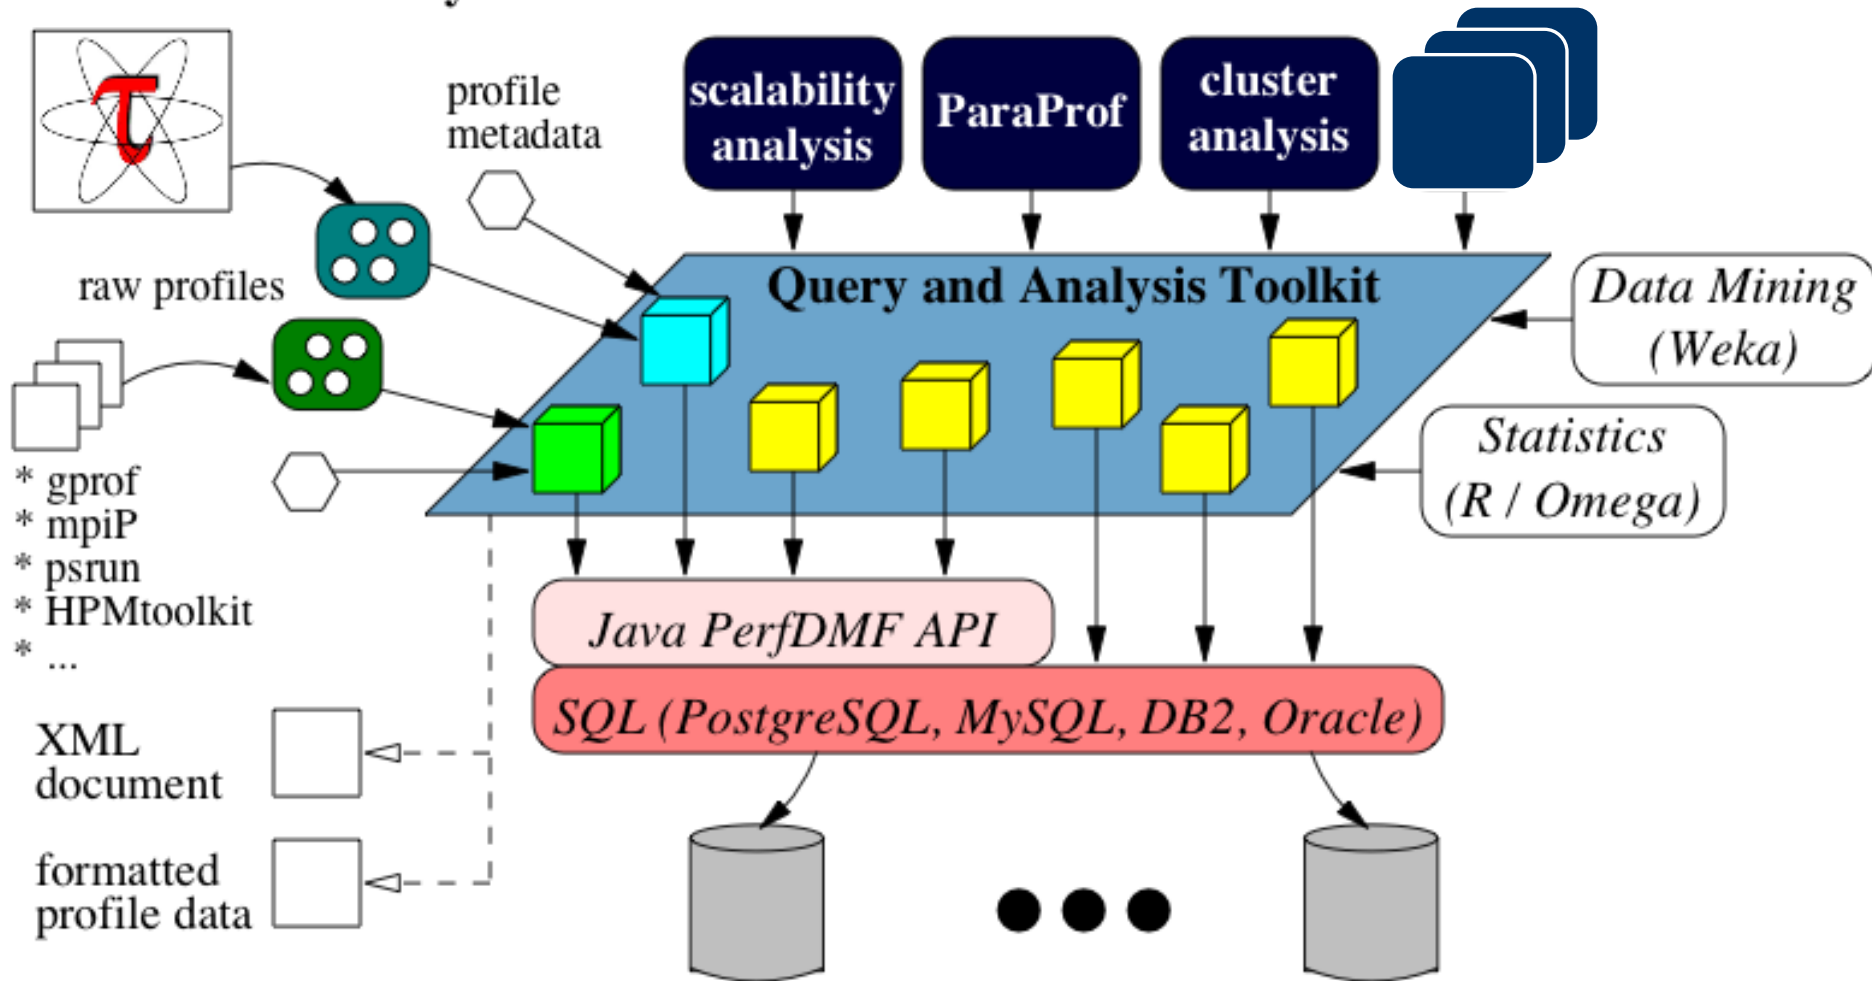
Performance Research Lab, University of Oregon

<http://TAU.uoregon.edu>



TAU Performance System

Performance Analysis Programs



- **Configure PerfDMF (Done by each user)**

- % perfdmf_configure --create-default

- Choose derby, PostgreSQL, MySQL, Oracle or DB2
 - Hostname
 - Username
 - Password
 - Say yes to downloading required drivers (we are not allowed to distribute these)
 - Stores parameters in your ~/.ParaProf/perfdmf.cfg file

- **Configure PerfExplorer (Done by each user)**

- % perfexplorer_configure

- **Execute PerfExplorer**

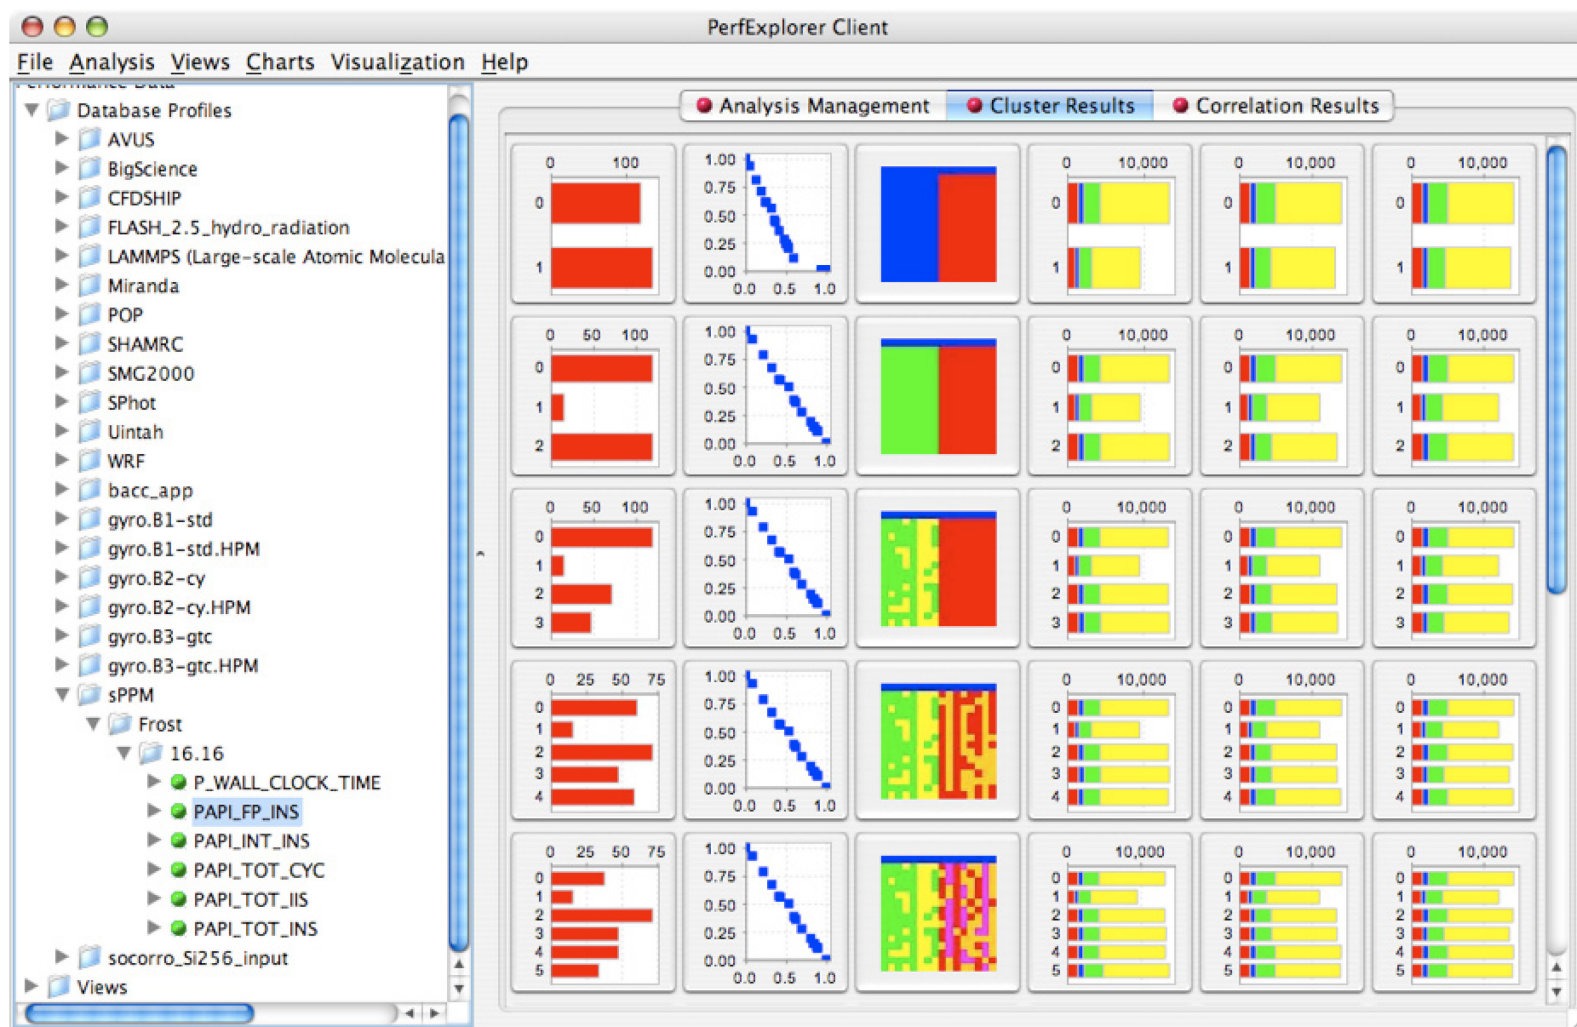
- % perfexplorer


```
% wget http://tau.uoregon.edu/data.tgz (Contains CUBE profiles from Score-P)
% perfdmf_configure --create-default
(Chooses derby, blank user/passwd, yes to save passwd, defaults)
% perfexplorer_configure
(Yes to load schema, defaults)
% paraprof
(load each trial: DB -> Add Trial -> Type (Paraprof Packed Profile) -> OK) OR use
    perfdmf_loadtrial -a "app" -x "experiment" -n "name" file.ppk
Then,
% perfexplorer
(Select experiment, Menu: Charts -> Speedup)
```

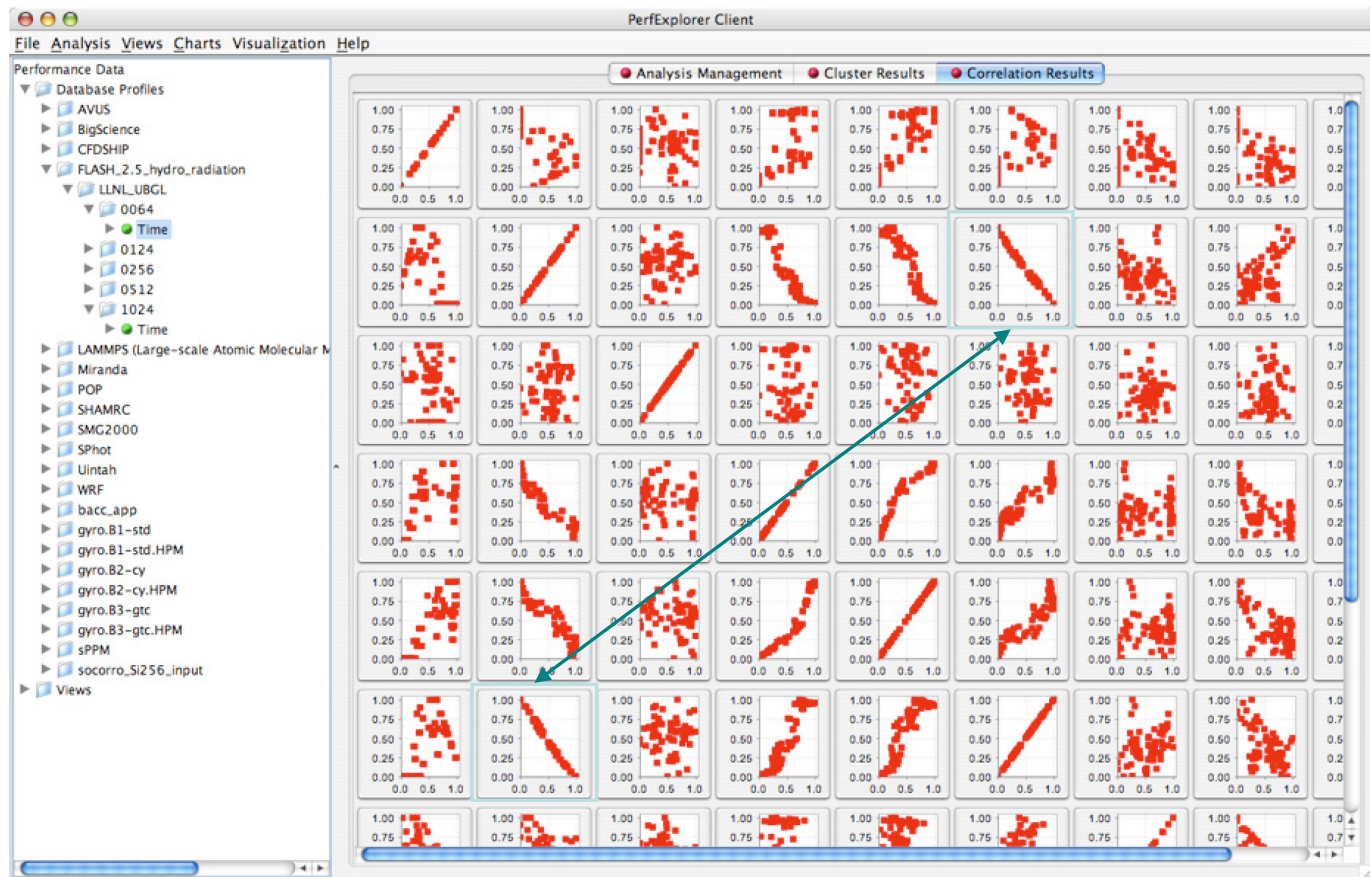
- Development of the TAU portal
 - Common repository for collaborative data sharing
 - Profile uploading, downloading, user management
 - Paraprof, PerfExplorer can be launched from the portal using Java Web Start (no TAU installation required)
- Portal URL
<http://tau.nic.uoregon.edu>

- Performance knowledge discovery framework
 - Data mining analysis applied to parallel performance data
 - comparative, clustering, correlation, dimension reduction, ...
 - Use the existing TAU infrastructure
 - TAU performance profiles, PerfDMF
 - Client-server based system architecture
- Technology integration
 - Java API and toolkit for portability
 - PerfDMF
 - R-project/Omegahat, Octave/Matlab statistical analysis
 - WEKA data mining package
 - JFreeChart for visualization, vector output (EPS, SVG)


- Performance data represented as vectors - each dimension is the cumulative time for an event
- *k*-means: *k* random centers are selected and instances are grouped with the "closest" (Euclidean) center
- New centers are calculated and the process repeated until stabilization or max iterations
- Dimension reduction necessary for meaningful results
- Virtual topology, summaries constructed



- Describes strength and direction of a linear relationship between two variables (events) in the data



- -0.995 indicates strong, negative relationship
- As CALC_CUT_BLOCK_CONTRIBUTIONS() increases in execution time, MPI_Barrier() decreases



- Relative speedup, efficiency
 - total runtime, by event, one event, by phase
- Breakdown of total runtime
- Group fraction of total runtime
- Correlating events to total runtime
- Timesteps per second

PerfExplorer Client

File Analysis Charts Help

▼ gyro.B1-std

- ▶ B1-std-hwpc.phoenix.0x002
- ▶ B1-std-inst.phoenix.0x002
- ▶ B1-std-inst.phoenix.0x002.profil
- ▶ B1-std-nl2.cheetah.affnosng
- ▶ B1-std-nl2.cheetah.affsng
- ▶ B1-std-nl2.cheetah.noaffnosng
- ▶ B1-std-nl2.phoenix.0x002
- ▶ B1-std-nl2.phoenix.0x002scr
- ▶ B1-std.53newest.phoenix.0x002
- ▶ B1-std.cheetah.affnosng
- ▶ B1-std.cheetah.affsng
- ▶ B1-std.cheetah.noaffnosng
- ▶ B1-std.hockney
- ▶ B1-std.new.phoenix.0x002
- ▶ B1-std.phoenix.0x002
- ▶ B1-std.phoenix.0x002scr
- ▶ B1-std.ram0x002.a
- ▶ B1-std.ram0x002.b
- ▼ B1-std.seaborg
 - ▶ B1-std.timing.seaborg.128
 - ▶ B1-std.timing.seaborg.16
 - ▶ B1-std.timing.seaborg.256
 - ▶ B1-std.timing.seaborg.32
 - ▶ B1-std.timing.seaborg.512
 - ▶ B1-std.timing.seaborg.64
- ▶ B1-std.tg
- ▶ gyro.B2-cy
- ▶ gyro.B3-gtc

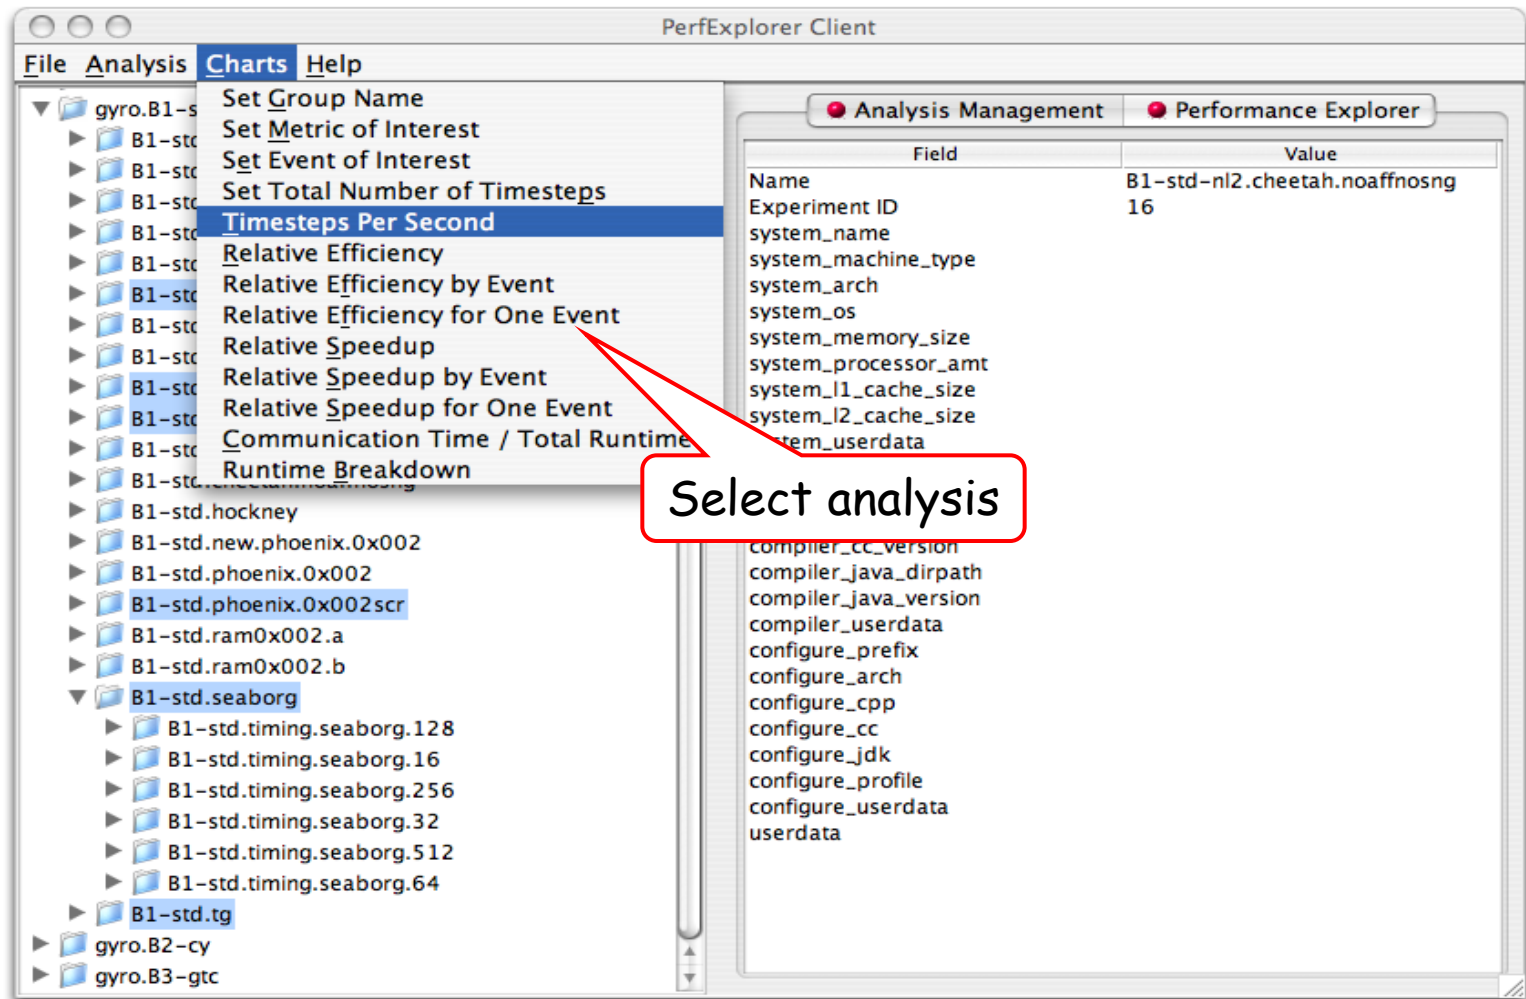
Analysis Management Performance Explorer


Field	Value
Name	B1-std-nl2.cheetah.noaffnosng
Experiment ID	16
system_name	
system_machine_type	
system_arch	
system_os	
system_memory_size	
system_processor_amt	
system_l1_cache_size	
system_l2_cache_size	
system_userdata	
compiler_cpp_name	
compiler_cpp_version	
configuration_prefix	
configuration_arch	


Experiment metadata


Select experiments and trials of interest


Data organized in application, experiment, trial structure (will allow arbitrary in future)















Usage Scenarios: Evaluate Scalability HPS

- Goal: How does my application scale? What bottlenecks occur at what core counts?
- Load profiles in PerfDMF database and examine with PerfExplorer



Usage Scenarios: Evaluate Scalability HPS








TAU/PerfExplorer: Total TIME Breakdown

Total TIME Breakdown for IRMHD:Scaling_BGP



■ ADVANCE_DIFFUSION
 ■ CCHEBYB
 ■ CCHEBYF
 ■ CHEBDIFF
 ■ DERIVE
 ■ ENERGY
 ■ MPIFFT::CRFFT2D_MPI
 ■ MPIFFT::HC2R
 ■ MPIFFT::R2HC
 ■ MPIFFT::RCFFT2D_MPI
 ■ MPIFFT::REORDER_COMPLEX
 ■ MPI_Allreduce()
 ■ MPI_Alltoall()
 ■ MPI_Barrier()
 ■ MPI_Init()
 ■ MPI_Waitall()
 ■ RK3NL::ERK3_STAGE
 ■ WDERIVS
 ■ other

Performance Regression Testing VI-HPS

