Metadata Collection for Performance Analysis

Karen L. Karavanic
Associate Professor of Computer Science
Portland State University

The PerfTrack Project

- PerfTrack is a tool for storing, exploring, and analyzing performance data
- Our Approach:
  - Collect and store as much information as possible about each build and run of an application
  - Integrate database technology into a performance analysis tool
  - Store a wide variety of performance data
    - Data from different measurement tools
      - Tracing, DPCL, Paradyn, TAU, Vampir, Speedshop, HW counters, etc.
    - Native application performance measurements

PerfTrack Design

- Design goal: No knowledge of database technology or vocabulary should be required to use PerfTrack
- Design goal: PerfTrack must be scalable to 1000s of program runs and 100,000s of performance results
- Design goal: PerfTrack must be flexible enough to store data from different measurement tools and different types of performance studies
- Design goal: PerfTrack must be extensible to accommodate future innovations in measurement and analysis
- Design goal: PerfTrack should not be limited to a specific DBMS package for its data store

PerfTrack Design

- Data Collection Scripts:
  - Build environment
  - Run environment
  - Performance Data
- PTdataStore interface:
  - Shelters PerfTrack user from the DBMS
  - Perftrack Data Format (PTDF)
- Data navigation and analysis
  - PerfTrack GUIs
  - Command line interface
  - Direct SQL query

PerfTrack Design

- Performance Result: a measured value
  - Metric: cpu time, wall clock time, wait time
  - Context: whole program? One function? One process?
  - Value: result of measuring a metric in a context
- Everything is a resource: machine, node, process, function, execution …
  - Resource attribute:
    - Execution - timestamp, build, machine, etc.
    - Compiler - version
  - Resource type
    - Hierarchical or non-hierarchical

Build Application

```
ptbuild.py --app umt2k  --pathToExe ./umt2k --srcDir . -V
```

This collects information about the build, such as the machine the application was built on, the compilers used, and environment variables that were set during the build.

build
- Time stamp
- Environment Variables

operatingSystem
- Name
- Release
- Version

machine
- Name

compiler
- Flags
- Include paths
- Libraries
- Path to compiler
- Vendor
- Version
- MPI script used
- MPI script flags
- Path to MPI script
- MPI script include paths
- MPI script libraries
- MPI script library paths

module
- Path to library
- Size
- Timestamp
- Type

Build Attributes

- BUILD: Date,Time,Environment Variables
- BUILD/MODULE: Path, Size,Timestamp
- COMPILER: CompileFlags,IncludePaths,Libraries,LibraryPaths,mpiScriptCompilerFlags,mpiScriptLibraries,mpiScriptLibraryPaths, mpiScriptName,mpiScriptPath,Path
- ENVIRONMENT/MODULE: LibraryDynamic, Path,Size,Timestamp,Type

Execute Application

```
ptrun.py --app umt2k  --batchFile psub.script --inputDeck “opacfile,rtin,smartin”
--exeName ./umt2k
```

This collects information about the execution, such as the machine the application was run on, and environment variables that were set during the execution.

environment

module
- Path to library
- Size
- Timestamp
- Type

operatingSystem
- Name
- Release
- Version

machine
- Name

node
- Name

execution
- App-specific values
- Permissions on exe
- Time stamp
- Number of processes
- Number of threads
- Concurrency model (MPI, OpenMP)
- Environment variables
- Languages in application
- User who ran application

processes
- Rank

threads
- Rank

inputDeck
- Name
- Mod time

submission
- Time stamp
- Name of batch file
- Commands in batch file
- Run command
- Environment variables set in batch file

Execute Application

performanceTool
- Name
- Version

metric
- Name

time
- Name
- Start Time
- End Time

interval
- Name
- Start Time
- End Time

subinterval
- Name
- Start Time
- End Time

performanceResults
- Metric
- Performance tool
- Value
- Units
- Start time
- End time
- Execution

Execution Attributes

- Concurrency
- Environment Variables
- Executable:
  - GID, Name, Permissions, Size, Timestamp, UID, j
- Job:
  - CompletionTime, Exit Status, Nodes, Resources Used, StartTime
- Languages
- Launch Date,Time
- NumberOfProcesses
- PageSize
- ProcessesPerNode, ThreadsPerProcess
- RunErrorMsg # any error messages from the job
- Username
- UsesMPI/OpenMP/Pthreads

PerfTrack Design: Generic Database Schema

resource_item

id	Integer
name	Varchar2(255)
type	Varchar2(255)
type_id	Integer
parent	Integer

resource_attribute

res_id	Integer
name	Varchar2(255)
value	Varchar2(255)
type	Varchar2(255)

resource_constraint

from	Integer
to	Integer

PerfTrack Design: Base Resource Types

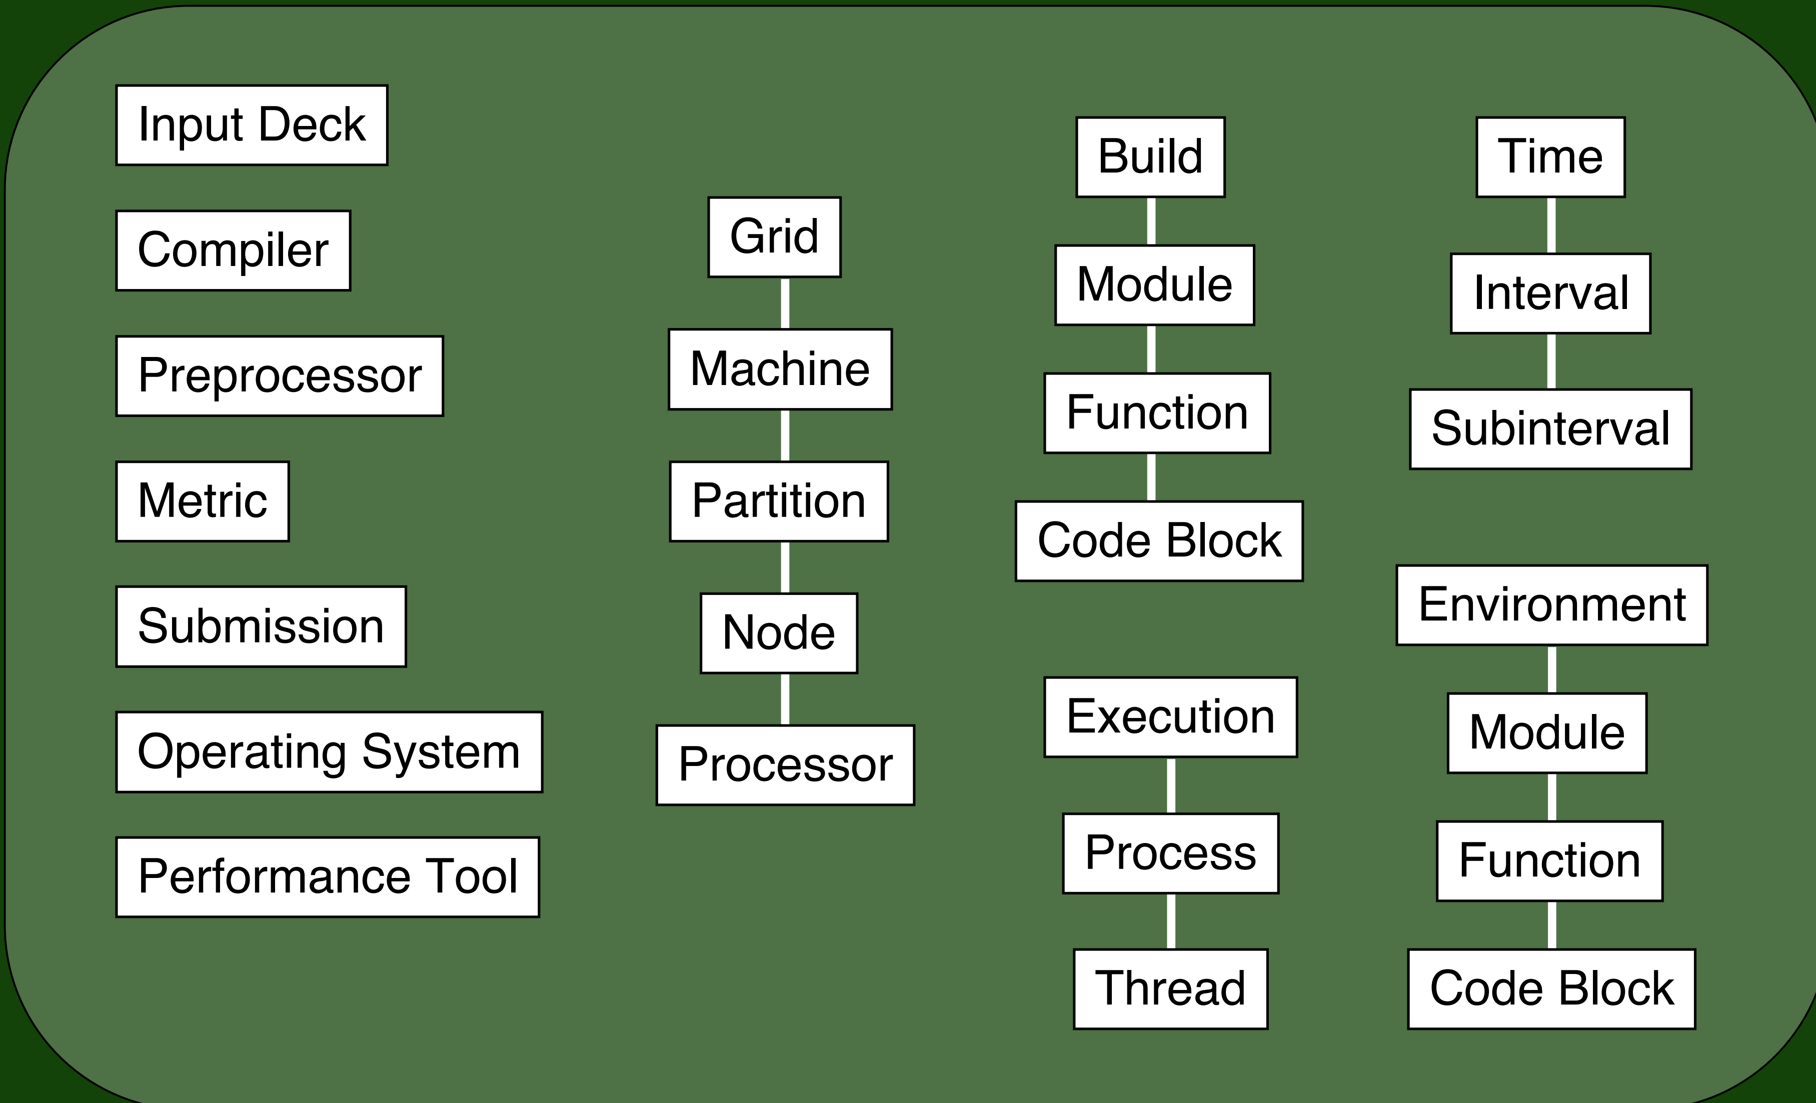

PerfTrack Design

- PerfTrack Data Format (PTdf):

```
ResourceType resourceTypeName

Application appName

Execution execName appName

Resource resourceName resourceTypeName execName
Resource resourceName resourceTypeName

ResourceAttribute resourceName attributeName
    attributeValue attributeType

ResourceConstraint resourceName1 resourceName2

PerfResult execName resourceSet perfToolName
    metricName value units startTime endTime
```

Flexible Schema: PERI project example

Base Resource Types

- Grid
- Build
- Environment
- Compiler
- Time
- Operating System
- Execution
- Input Deck
- Preprocessor
- Performance Tool
- Submission
- Metric

Custom resource types

- File System
- FileSystem/device

Custom Attributes

- Submission: Batch queue Entry
- Submission: PBS resources

Submission Attributes

- batchCmd # commands in the batch file
- batchFile
  - batchFileDateTime
- batchQueueEntry # entries in the queue at the time of submit
- launcher
  - launcherVersion
- machinePartition
- PBS Resources
- runCmd # run commands in the batch file

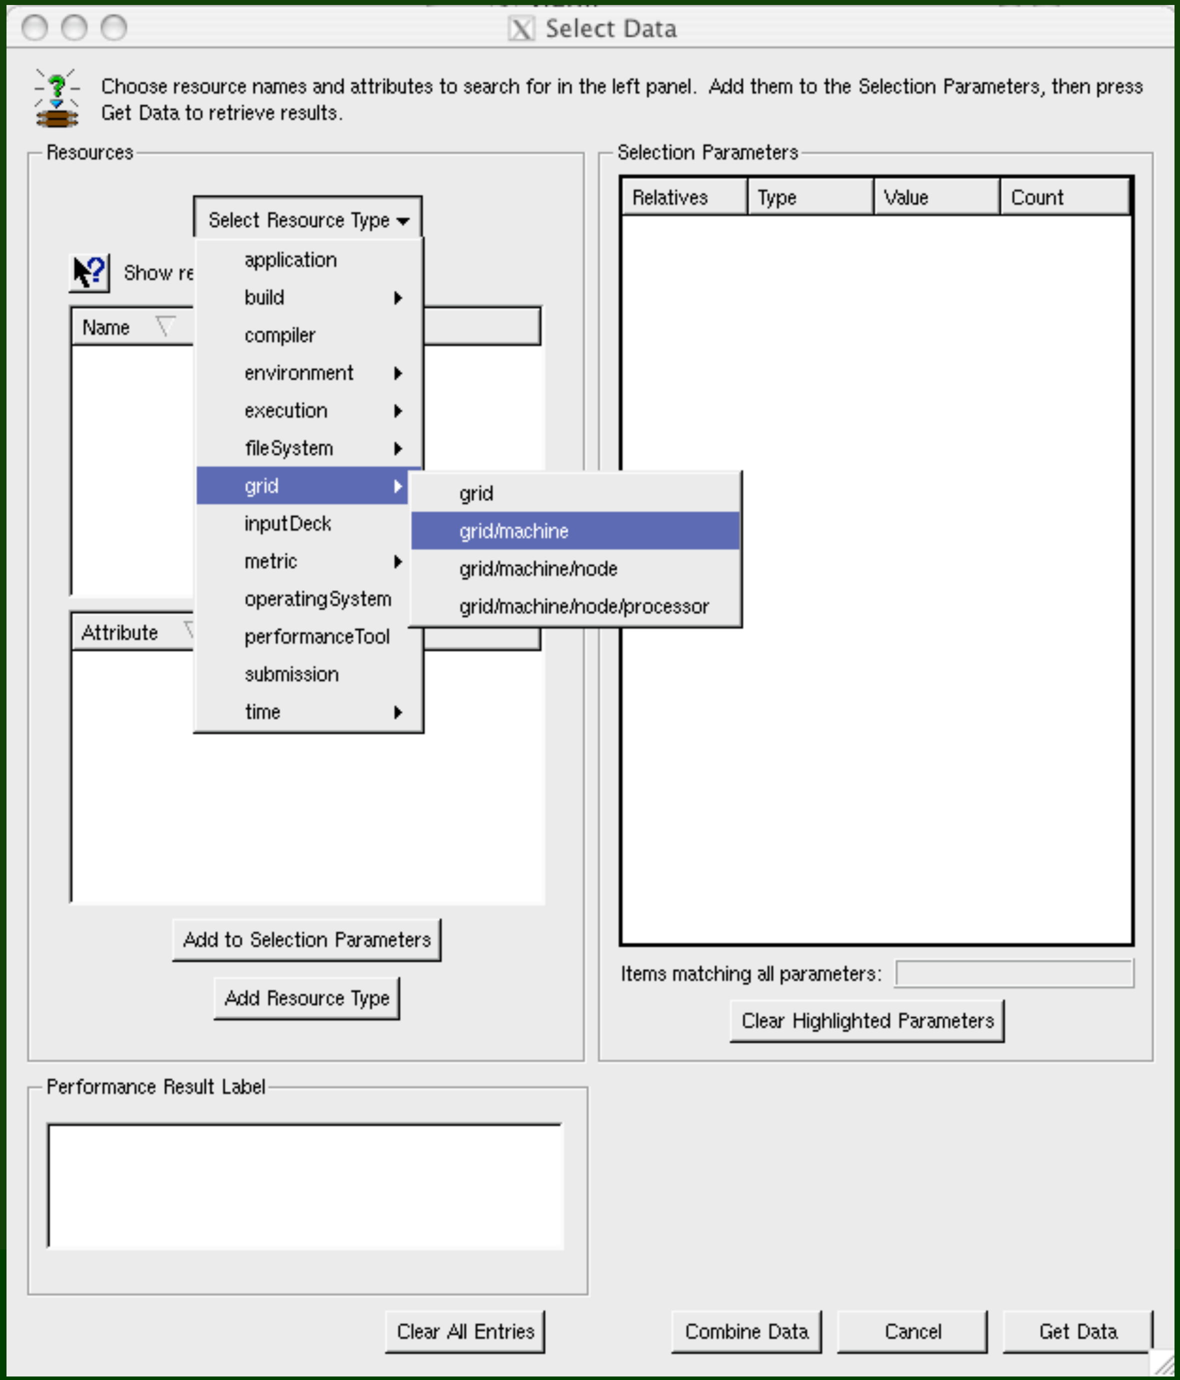
Select Data
Choose resource names and attributes to search for in the left panel. Add them to the Selection Parameters, then press Get Data to retrieve results.
Resources
Select Resource Type
Show re
Name
Attribute
application
build
compiler
environment
execution
fileSystem
grid
inputDeck
metric
operatingSystem
performanceTool
submission
time
grid
grid/machine
grid/machine/node
grid/machine/node/processor
Add to Selection Parameters
Add Resource Type
Selection Parameters
Relatives
Type
Value
Count
Items matching all parameters:
Clear Highlighted Parameters
Performance Result Label
Clear All Entries
Combine Data
Cancel
Get Data

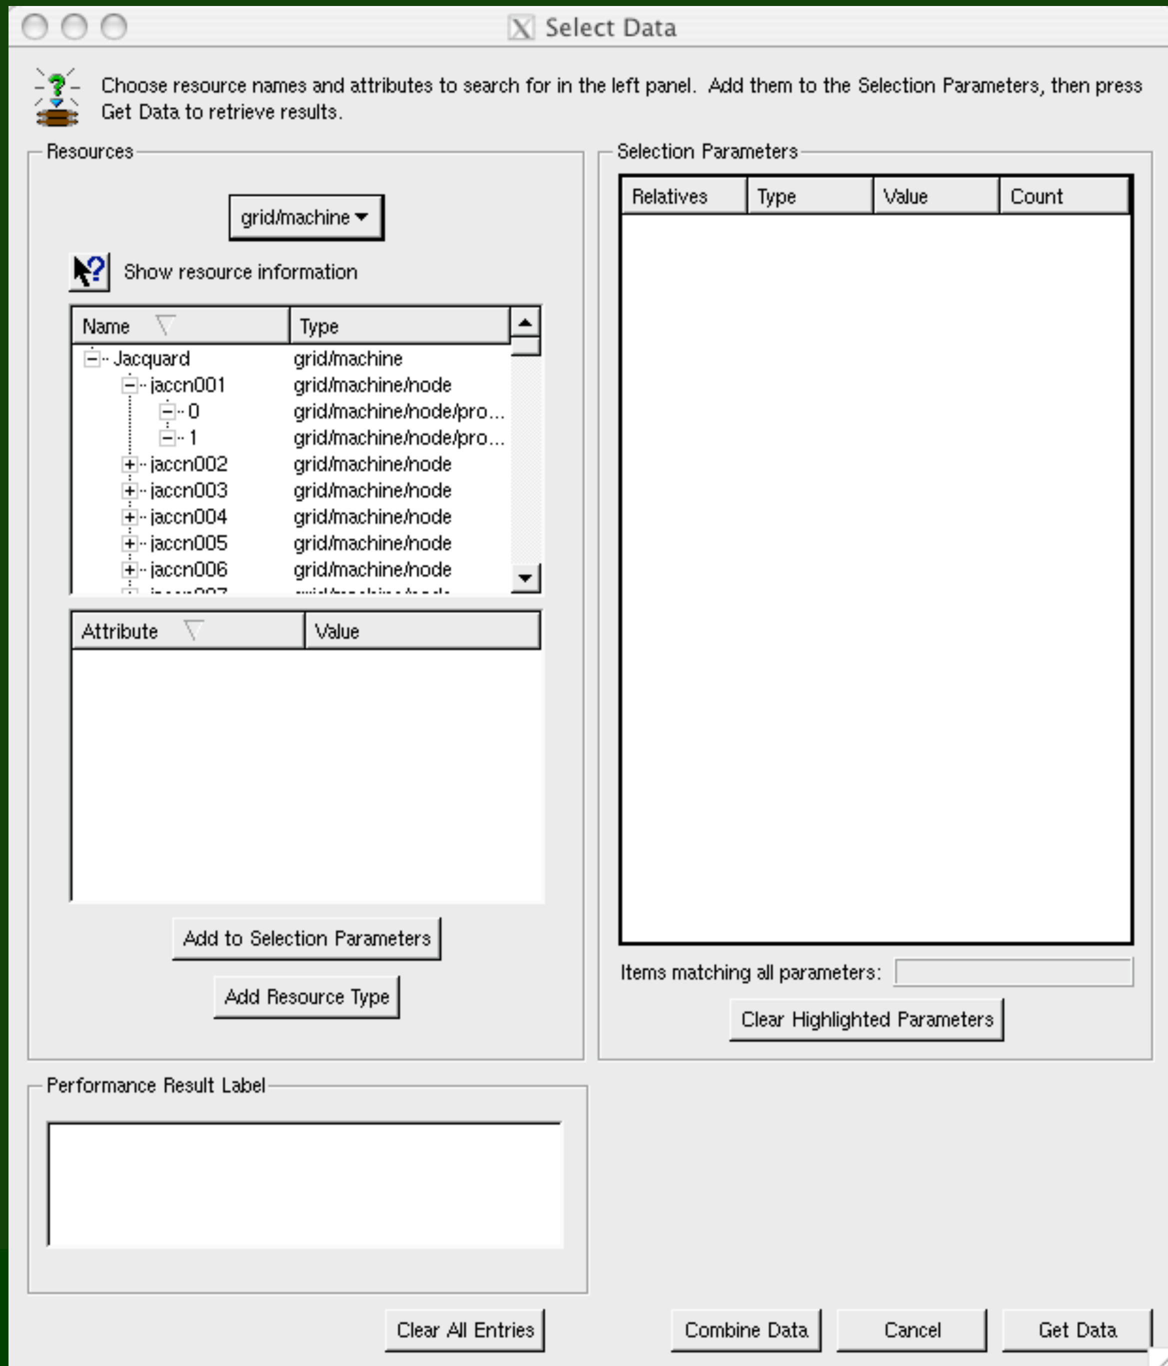
Select Data
Choose resource names and attributes to search for in the left panel. Add them to the Selection Parameters, then press Get Data to retrieve results.
Resources
grid/machine
Show resource information
Name
Type
Jacquard
grid/machine
jaccn001
grid/machine/node
0
grid/machine/node/pro...
1
grid/machine/node/pro...
jaccn002
grid/machine/node
jaccn003
grid/machine/node
jaccn004
grid/machine/node
jaccn005
grid/machine/node
jaccn006
grid/machine/node
Attribute
Value
Add to Selection Parameters
Add Resource Type
Selection Parameters
Relatives
Type
Value
Count
Items matching all parameters:
Clear Highlighted Parameters
Performance Result Label
Clear All Entries
Combine Data
Cancel
Get Data

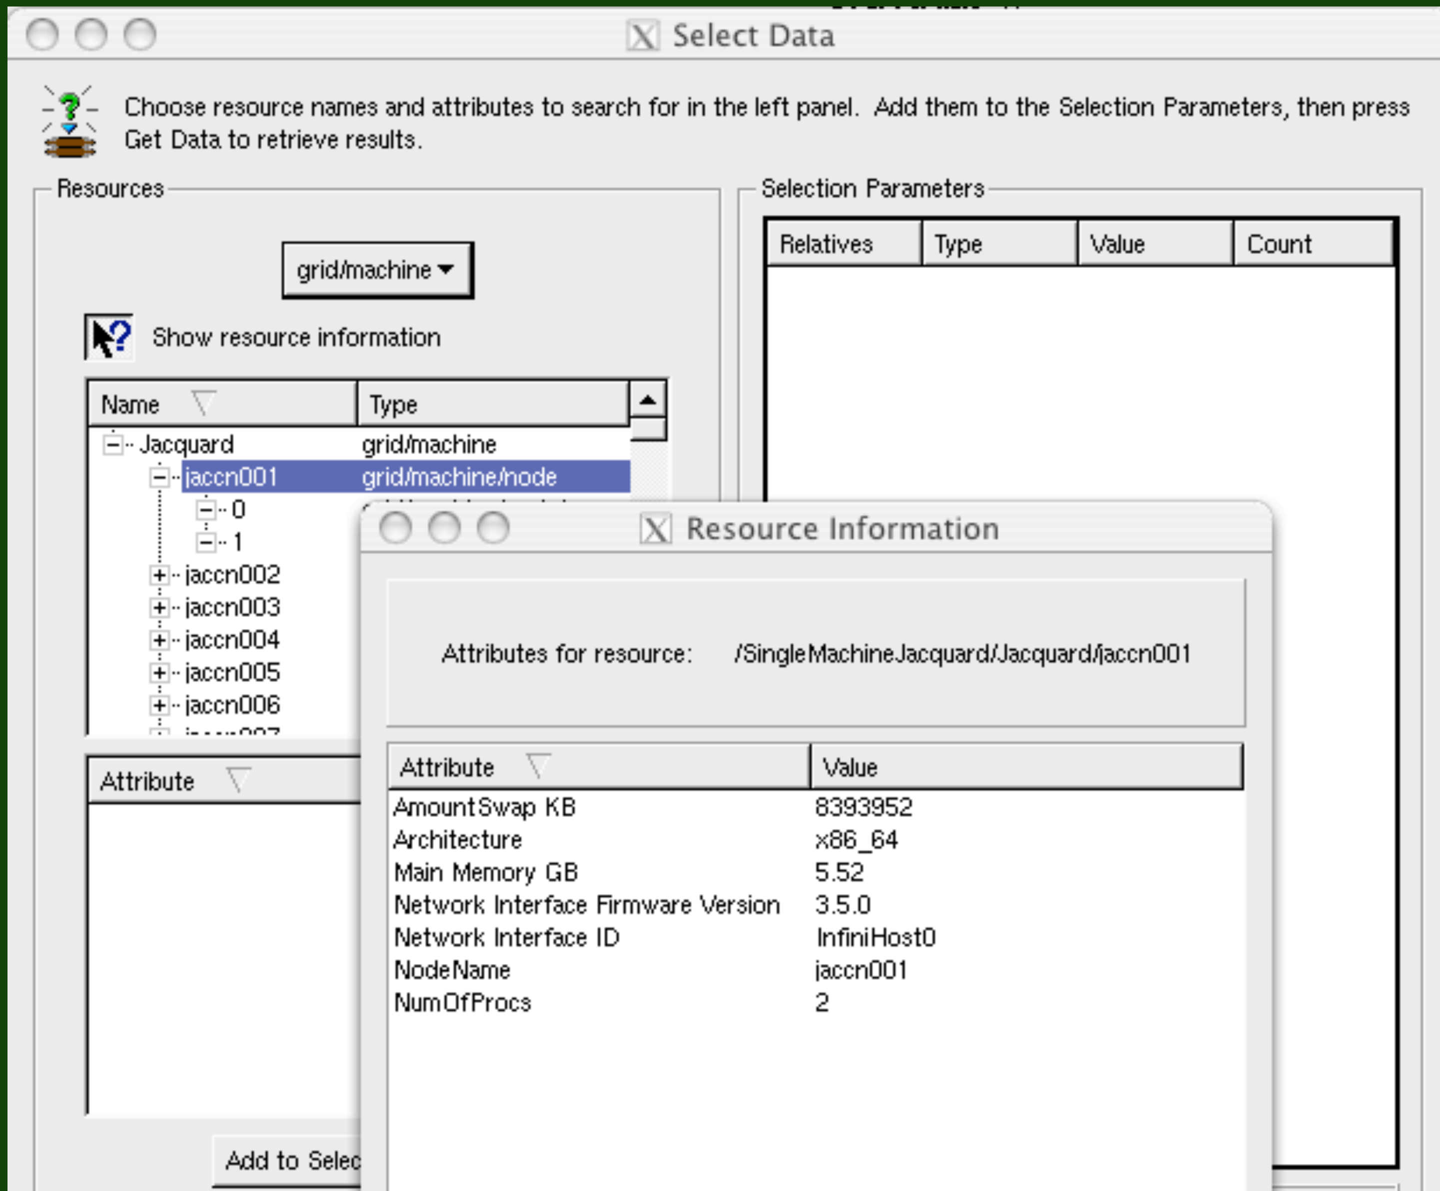
Select Data
Choose resource names and attributes to search for in the left panel. Add them to the Selection Parameters, then press Get Data to retrieve results.
Resources
grid/machine
Show resource information
Name
Type
Jacquard
grid/machine
jaccn001
grid/machine/node
0
1
jaccn002
jaccn003
jaccn004
jaccn005
jaccn006
Attribute
Add to Selec
Selection Parameters
Relatives
Type
Value
Count
Resource Information
Attributes for resource: /SingleMachineJacquard/Jacquard/jaccn001
Attribute
Value
AmountSwap KB
8393952
Architecture
x86_64
Main Memory GB
5.52
Network Interface Firmware Version
3.5.0
Network Interface ID
InfiniHost0
NodeName
jaccn001
NumOfProcs
2

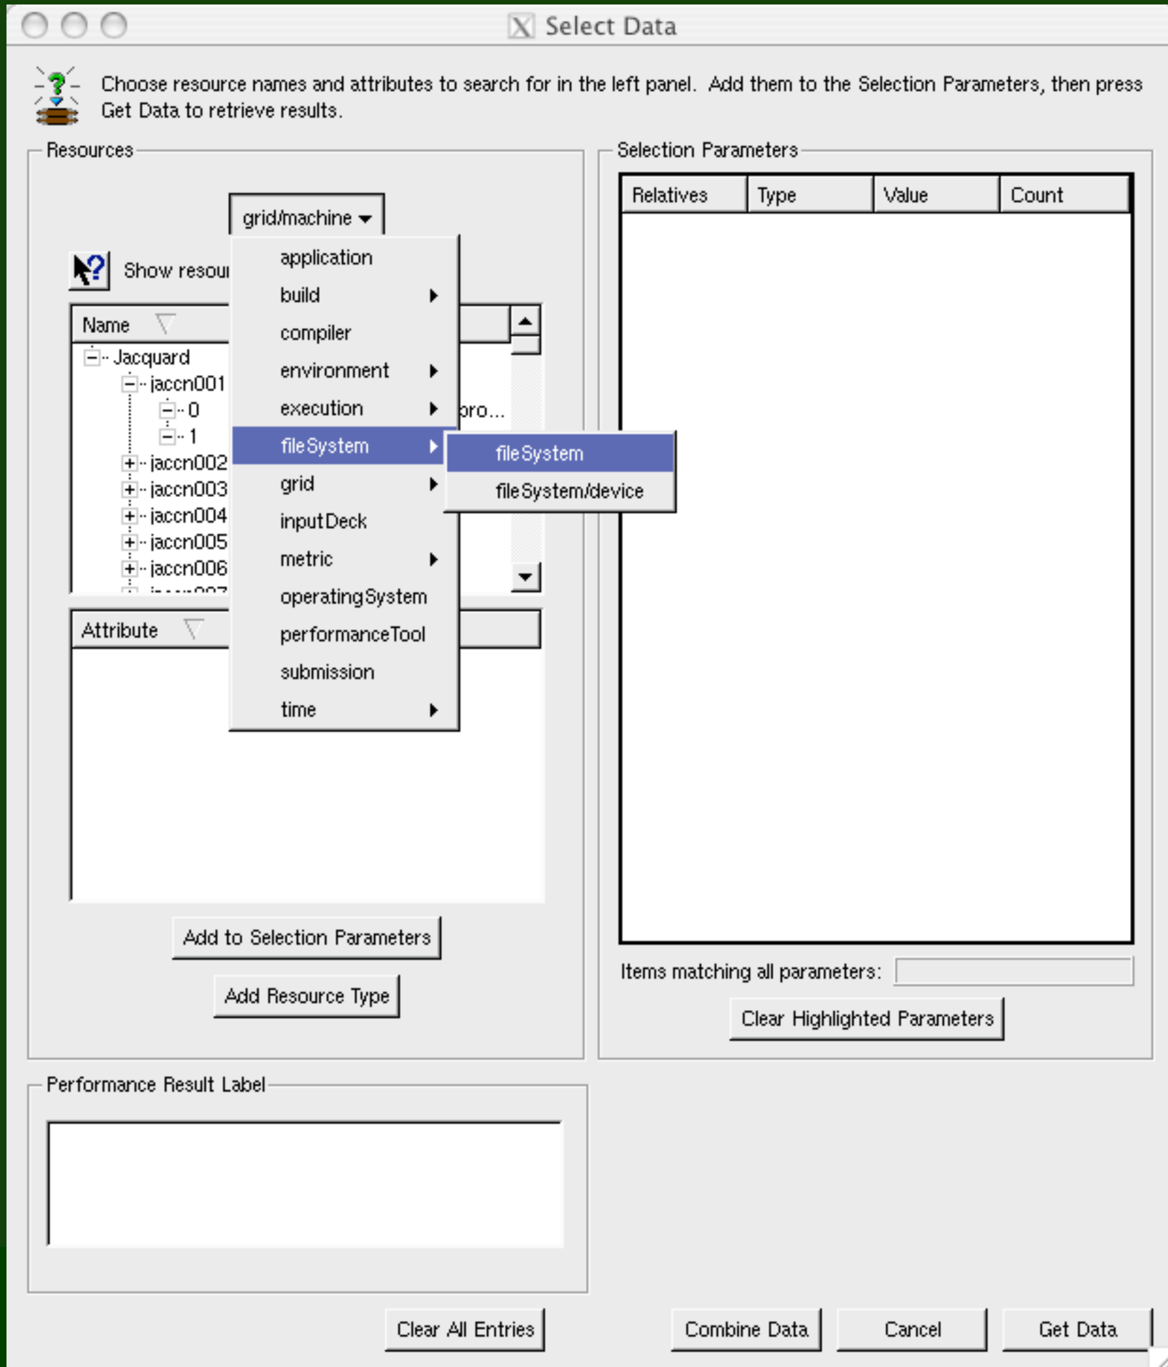
Select Data
Choose resource names and attributes to search for in the left panel. Add them to the Selection Parameters, then press Get Data to retrieve results.
Resources
grid/machine
Show resou
Name
Jacquard
jaccn001
0
1
jaccn002
jaccn003
jaccn004
jaccn005
jaccn006
pro...
application
build
compiler
environment
execution
fileSystem
grid
inputDeck
metric
operatingSystem
performanceTool
submission
time
fileSystem
fileSystem/device
Attribute
Add to Selection Parameters
Add Resource Type
Selection Parameters
Relatives
Type
Value
Count
Items matching all parameters:
Clear Highlighted Parameters
Performance Result Label
Clear All Entries
Combine Data
Cancel
Get Data

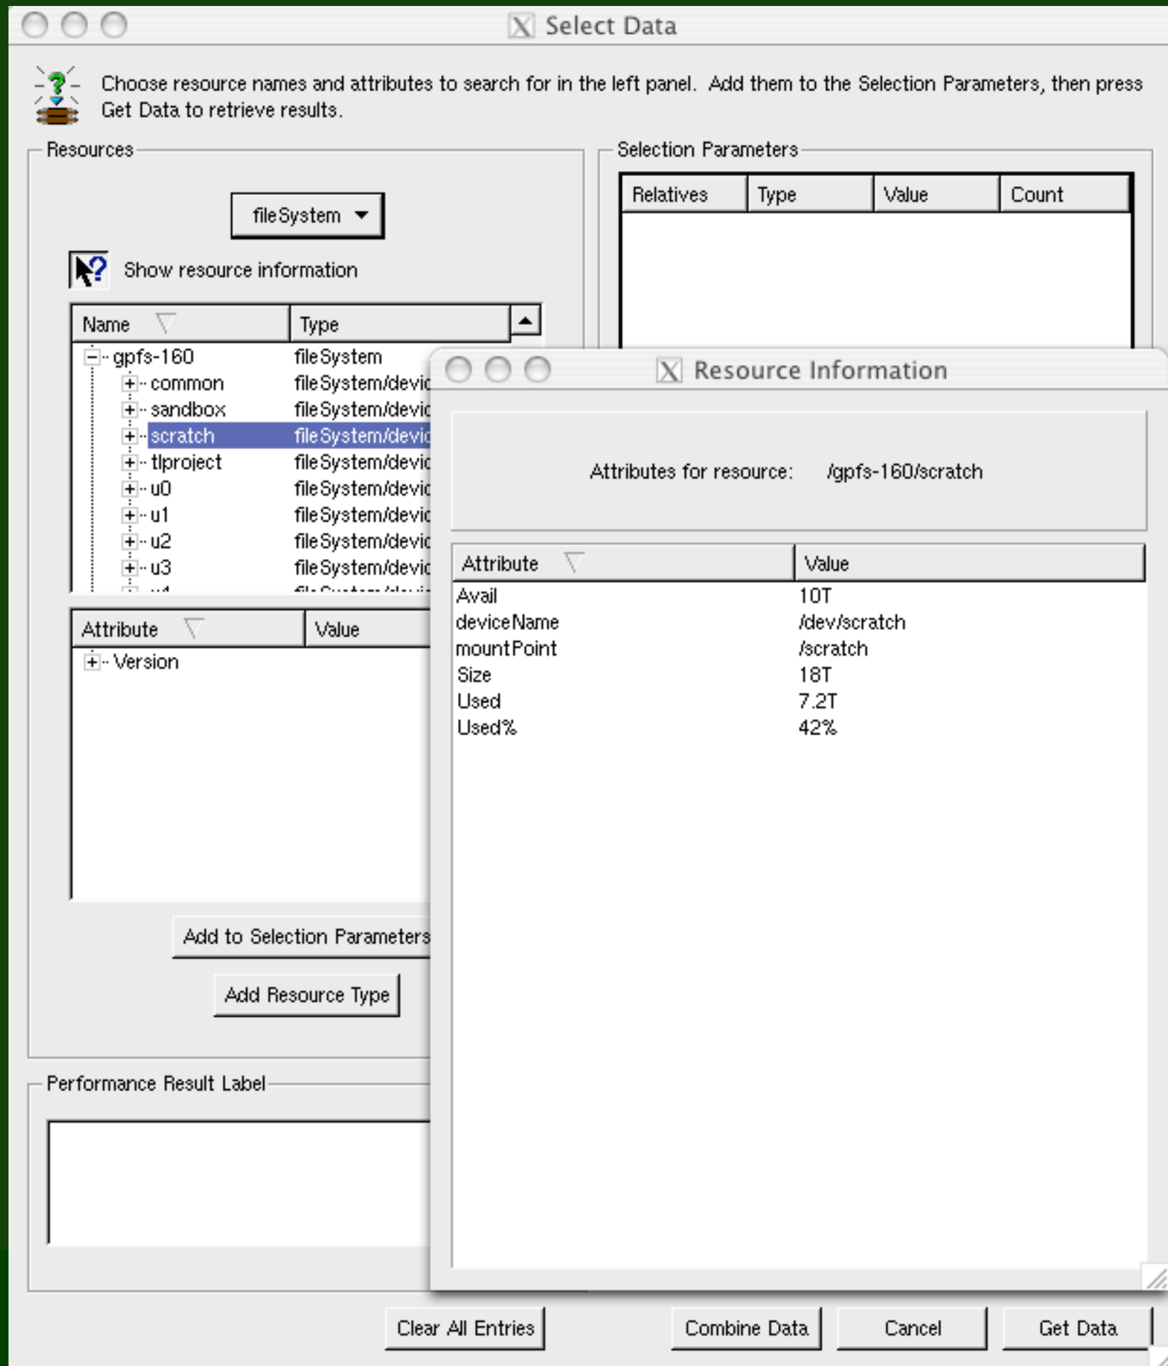
Select Data
Choose resource names and attributes to search for in the left panel. Add them to the Selection Parameters, then press Get Data to retrieve results.
Resources
fileSystem
Show resource information
Name
Type
gpfs-160
fileSystem
common
sandbox
scratch
tlproject
u0
u1
u2
u3
Attribute
Value
Version
Add to Selection Parameters
Add Resource Type
Performance Result Label
Selection Parameters
Relatives
Type
Value
Count
Resource Information
Attributes for resource: /gpfs-160/scratch
Attribute
Value
Avail 10T
deviceName /dev/scratch
mountPoint /scratch
Size 18T
Used 7.2T
Used% 42%
Clear All Entries
Combine Data
Cancel
Get Data

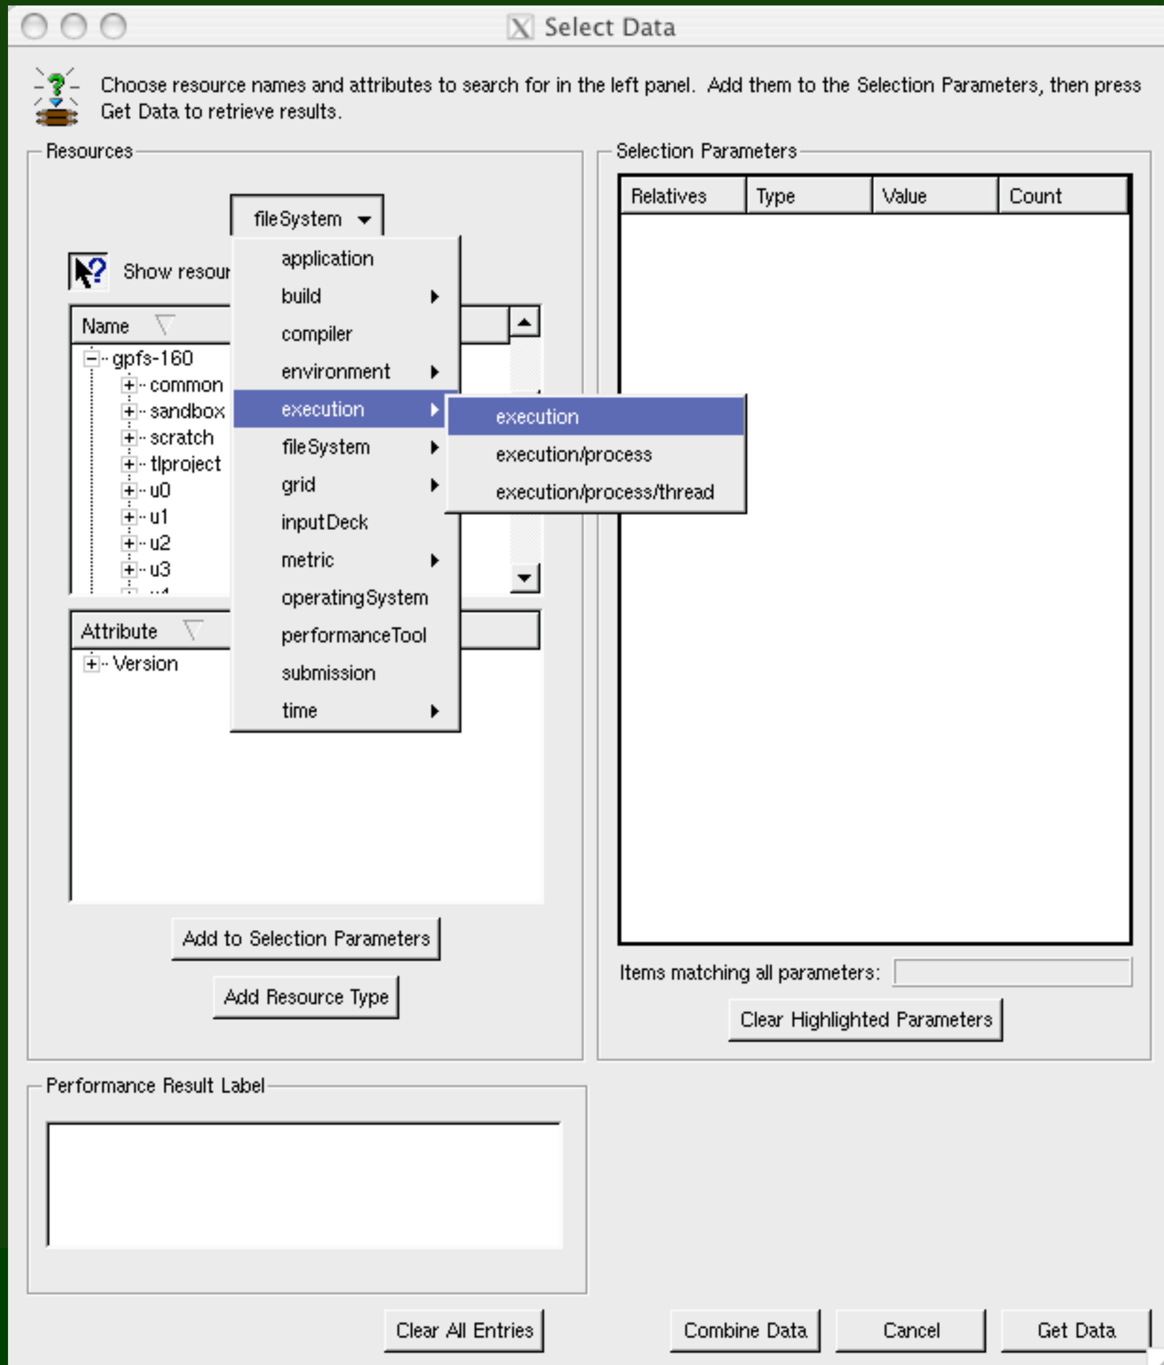
Select Data
Choose resource names and attributes to search for in the left panel. Add them to the Selection Parameters, then press Get Data to retrieve results.
Resources
fileSystem
Show resour
Name
gpfs-160
common
sandbox
scratch
tlproject
u0
u1
u2
u3
application
build
compiler
environment
execution
fileSystem
grid
inputDeck
metric
operatingSystem
performanceTool
submission
time
execution
execution/process
execution/process/thread
Attribute
Version
Add to Selection Parameters
Add Resource Type
Selection Parameters
Relatives
Type
Value
Count
Items matching all parameters:
Clear Highlighted Parameters
Performance Result Label
Clear All Entries
Combine Data
Cancel
Get Data

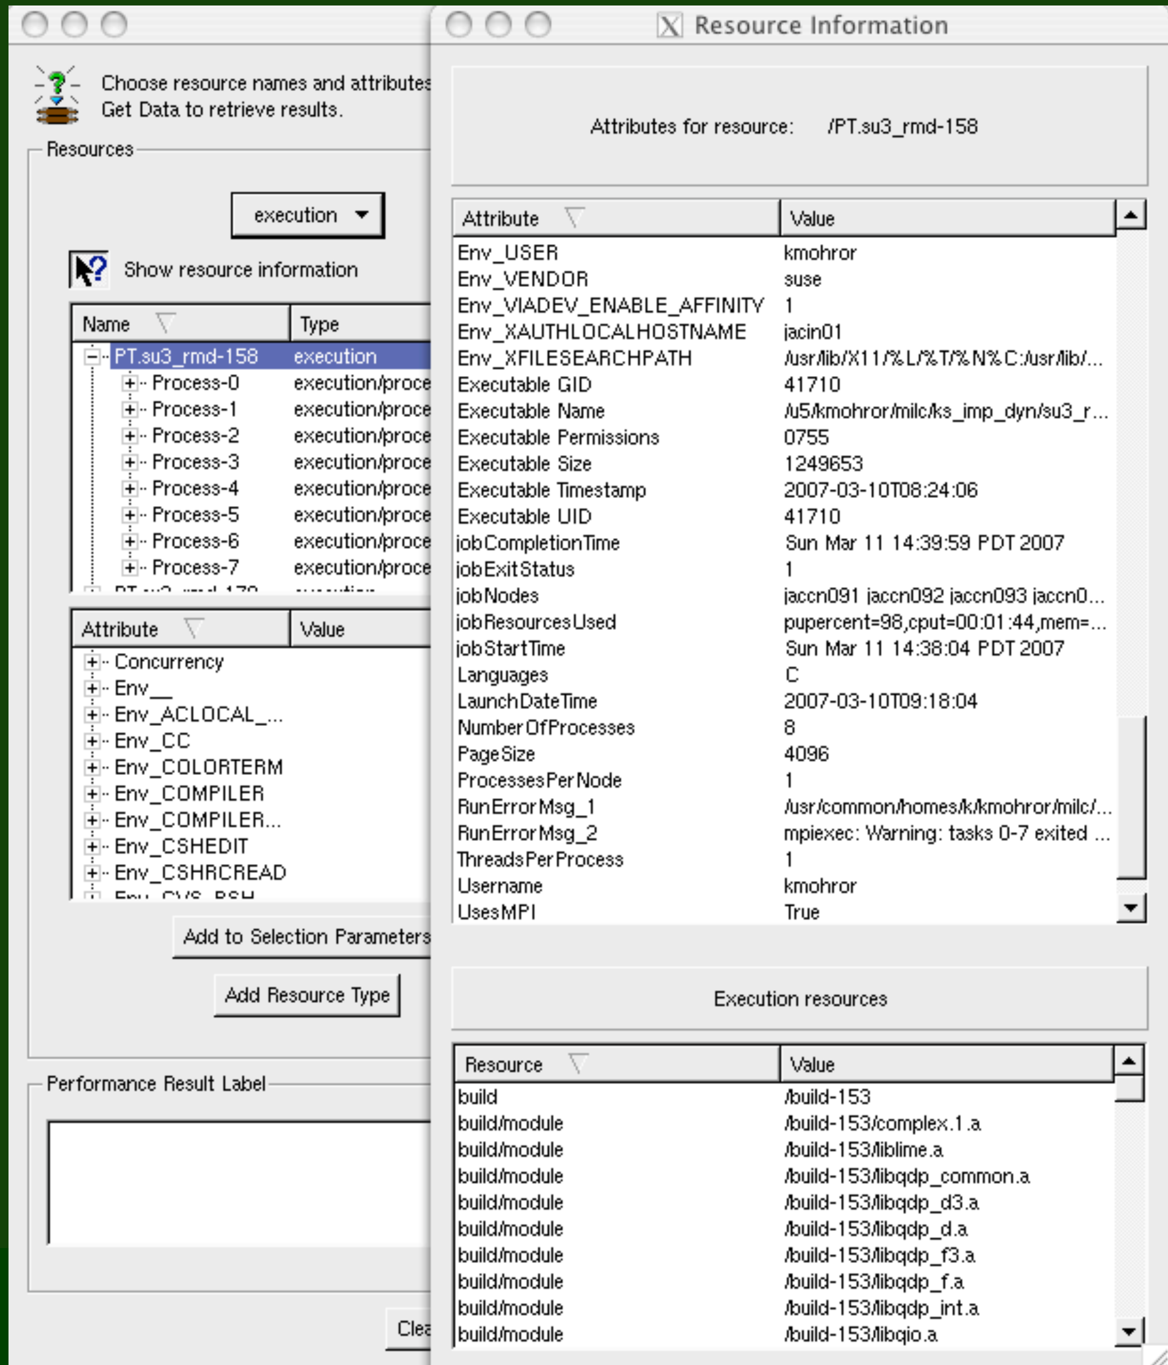
Choose resource names and attributes
Get Data to retrieve results.
Resources
execution
Show resource information
Name
Type
PT.su3_rmd-158
execution
Process-0
Process-1
Process-2
Process-3
Process-4
Process-5
Process-6
Process-7
execution/proce
Attribute
Value
Concurrency
Env__
Env_ACLOCAL_...
Env_CC
Env_COLORTERM
Env_COMPILER
Env_COMPILER...
Env_CSHEDIT
Env_CSHRCREAD
Add to Selection Parameters
Add Resource Type
Performance Result Label
Resource Information
Attributes for resource: /PT.su3_rmd-158
Attribute
Value
Env_USER kmohror
Env_VENDOR suse
Env_VIADEV_ENABLE_AFFINITY 1
Env_XAUTHLOCALHOSTNAME jacin01
Env_XFILESEARCHPATH /usr/lib/X11/%L/%T/%N%C:/usr/lib/...
Executable GID 41710
Executable Name /u5/kmohror/milc/ks_imp_dyn/su3_r...
Executable Permissions 0755
Executable Size 1249653
Executable Timestamp 2007-03-10T08:24:06
Executable UID 41710
jobCompletionTime Sun Mar 11 14:39:59 PDT 2007
jobExitStatus 1
jobNodes jaccn091 jaccn092 jaccn093 jaccn0...
jobResourcesUsed pupercent=98,cput=00:01:44,mem=...
jobStartTime Sun Mar 11 14:38:04 PDT 2007
Languages C
LaunchDateTime 2007-03-10T09:18:04
NumberOfProcesses 8
PageSize 4096
ProcessesPerNode 1
RunErrorMsg_1 /usr/common/homes/k/kmohror/milc/...
RunErrorMsg_2 mpiexec: Warning: tasks 0-7 exited ...
ThreadsPerProcess 1
Username kmohror
UsesMPI True
Execution resources
Resource
Value
build /build-153
build/module /build-153/complex.1.a
build/module /build-153/liblime.a
build/module /build-153/libqdp_common.a
build/module /build-153/libqdp_d3.a
build/module /build-153/libqdp_d.a
build/module /build-153/libqdp_f3.a
build/module /build-153/libqdp_f.a
build/module /build-153/libqdp_int.a
build/module /build-153/libqio.a

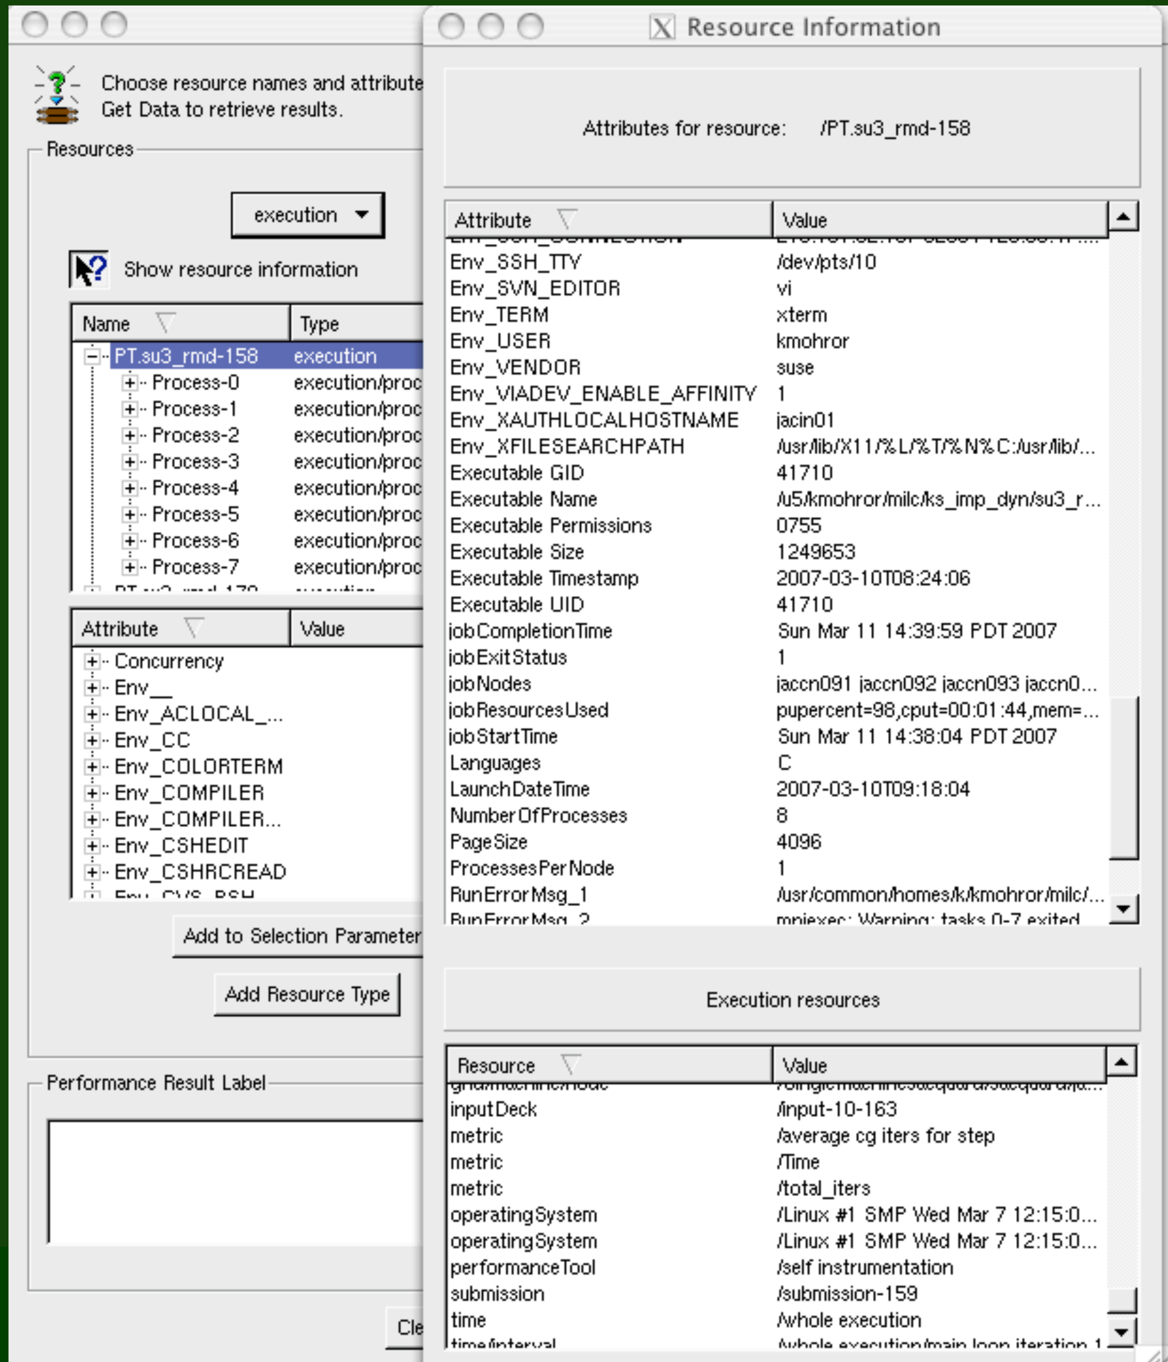
Resource Information
Choose resource names and attribute
Get Data to retrieve results.
Resources
execution
Show resource information
Name
Type
PT.su3_rmd-158 execution
Process-0 execution/proc
Process-1 execution/proc
Process-2 execution/proc
Process-3 execution/proc
Process-4 execution/proc
Process-5 execution/proc
Process-6 execution/proc
Process-7 execution/proc
Attribute
Value
Concurrency
Env__
Env_ACLOCAL_...
Env_CC
Env_COLORTERM
Env_COMPILER
Env_COMPILER...
Env_CSHEDIT
Env_CSHRCREAD
Add to Selection Parameter
Add Resource Type
Performance Result Label
Attributes for resource: /PT.su3_rmd-158
Attribute
Value
Env_SSH_TTY /dev/pts/10
Env_SVN_EDITOR vi
Env_TERM xterm
Env_USER kmohror
Env_VENDOR suse
Env_VIADEV_ENABLE_AFFINITY 1
Env_XAUTHLOCALHOSTNAME jacin01
Env_XFILESEARCHPATH /usr/lib/X11/%L/%T/%N%C:/usr/lib/...
Executable GID 41710
Executable Name /u5/kmohror/milc/ks_imp_dyn/su3_r...
Executable Permissions 0755
Executable Size 1249653
Executable Timestamp 2007-03-10T08:24:06
Executable UID 41710
jobCompletionTime Sun Mar 11 14:39:59 PDT 2007
jobExitStatus 1
jobNodes jaccn091 jaccn092 jaccn093 jaccn0...
jobResourcesUsed pupercent=98,cput=00:01:44,mem=...
jobStartTime Sun Mar 11 14:38:04 PDT 2007
Languages C
LaunchDateTime 2007-03-10T09:18:04
NumberOfProcesses 8
PageSize 4096
ProcessesPerNode 1
RunErrorMsg_1 /usr/common/homes/k/kmohror/milc/...
Execution resources
Resource
Value
inputDeck /input-10-163
metric /average cg iters for step
metric /Time
metric /total_iters
operatingSystem /Linux #1 SMP Wed Mar 7 12:15:0...
operatingSystem /Linux #1 SMP Wed Mar 7 12:15:0...
performanceTool /self instrumentation
submission /submission-159
time /whole execution

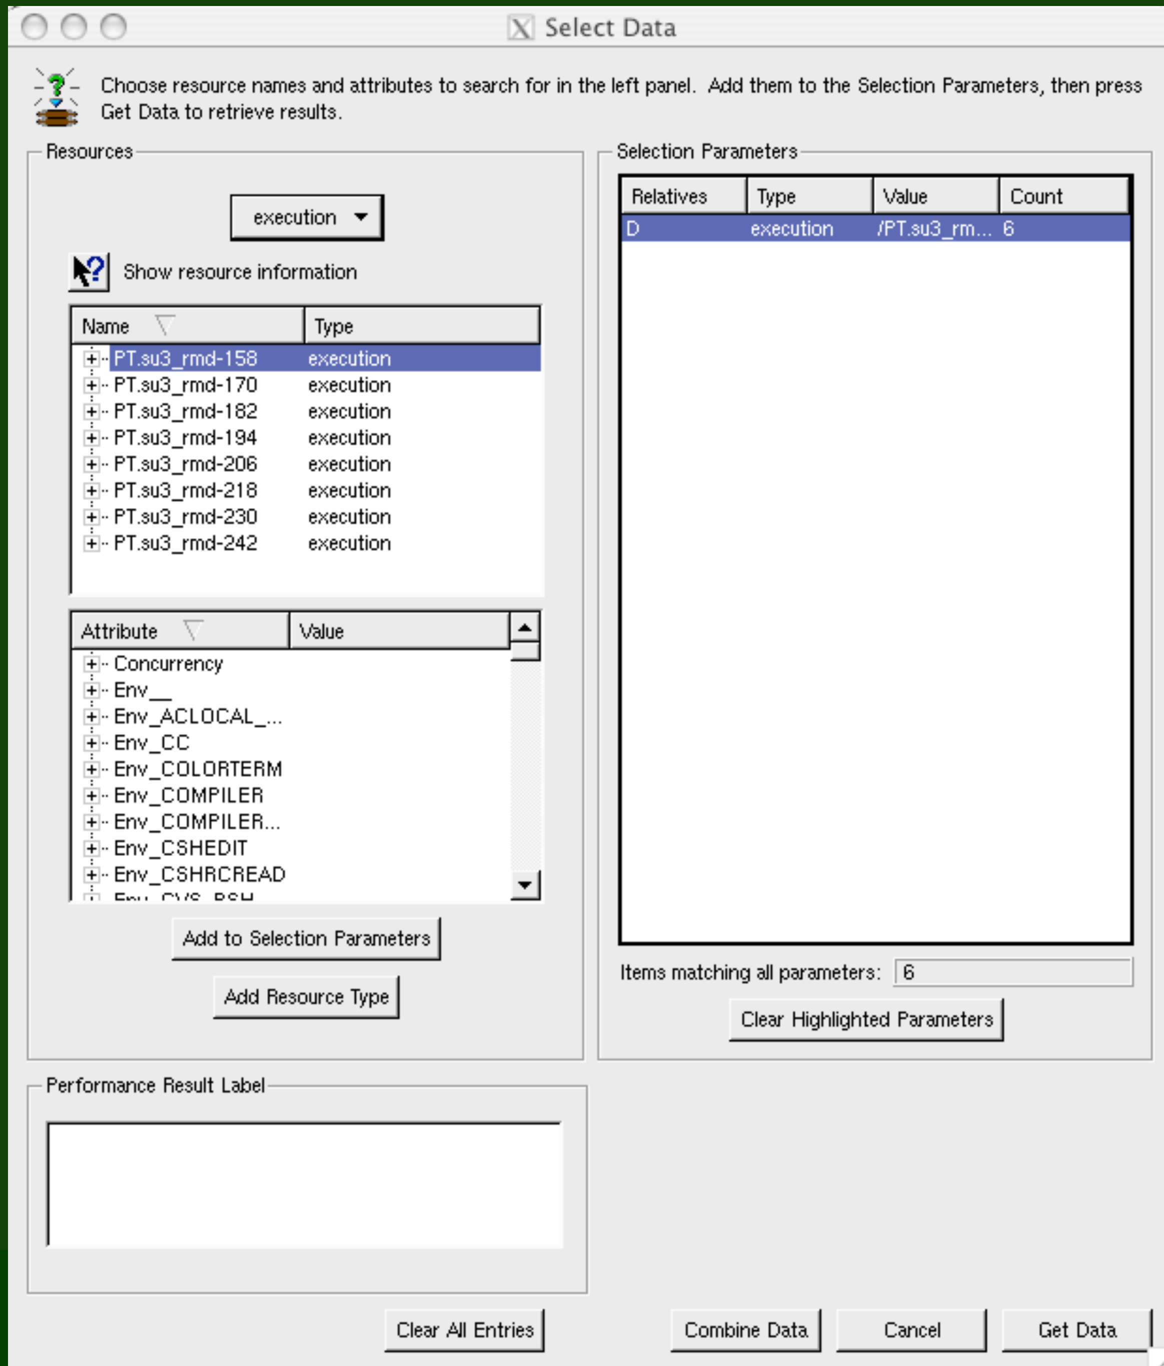

Select Data
Choose resource names and attributes to search for in the left panel. Add them to the Selection Parameters, then press Get Data to retrieve results.
Resources
execution
Show resource information
Name
Type
PT.su3_rmd-158 execution
PT.su3_rmd-170 execution
PT.su3_rmd-182 execution
PT.su3_rmd-194 execution
PT.su3_rmd-206 execution
PT.su3_rmd-218 execution
PT.su3_rmd-230 execution
PT.su3_rmd-242 execution
Attribute
Value
Concurrency
Env__
Env_ACLOCAL_...
Env_CC
Env_COLORTERM
Env_COMPILER
Env_COMPILER...
Env_CSHEDIT
Env_CSHRCREAD
Add to Selection Parameters
Add Resource Type
Selection Parameters
Relatives
Type
Value
Count
D execution /PT.su3_rm... 6
Items matching all parameters: 6
Clear Highlighted Parameters
Performance Result Label
Clear All Entries
Combine Data
Cancel
Get Data

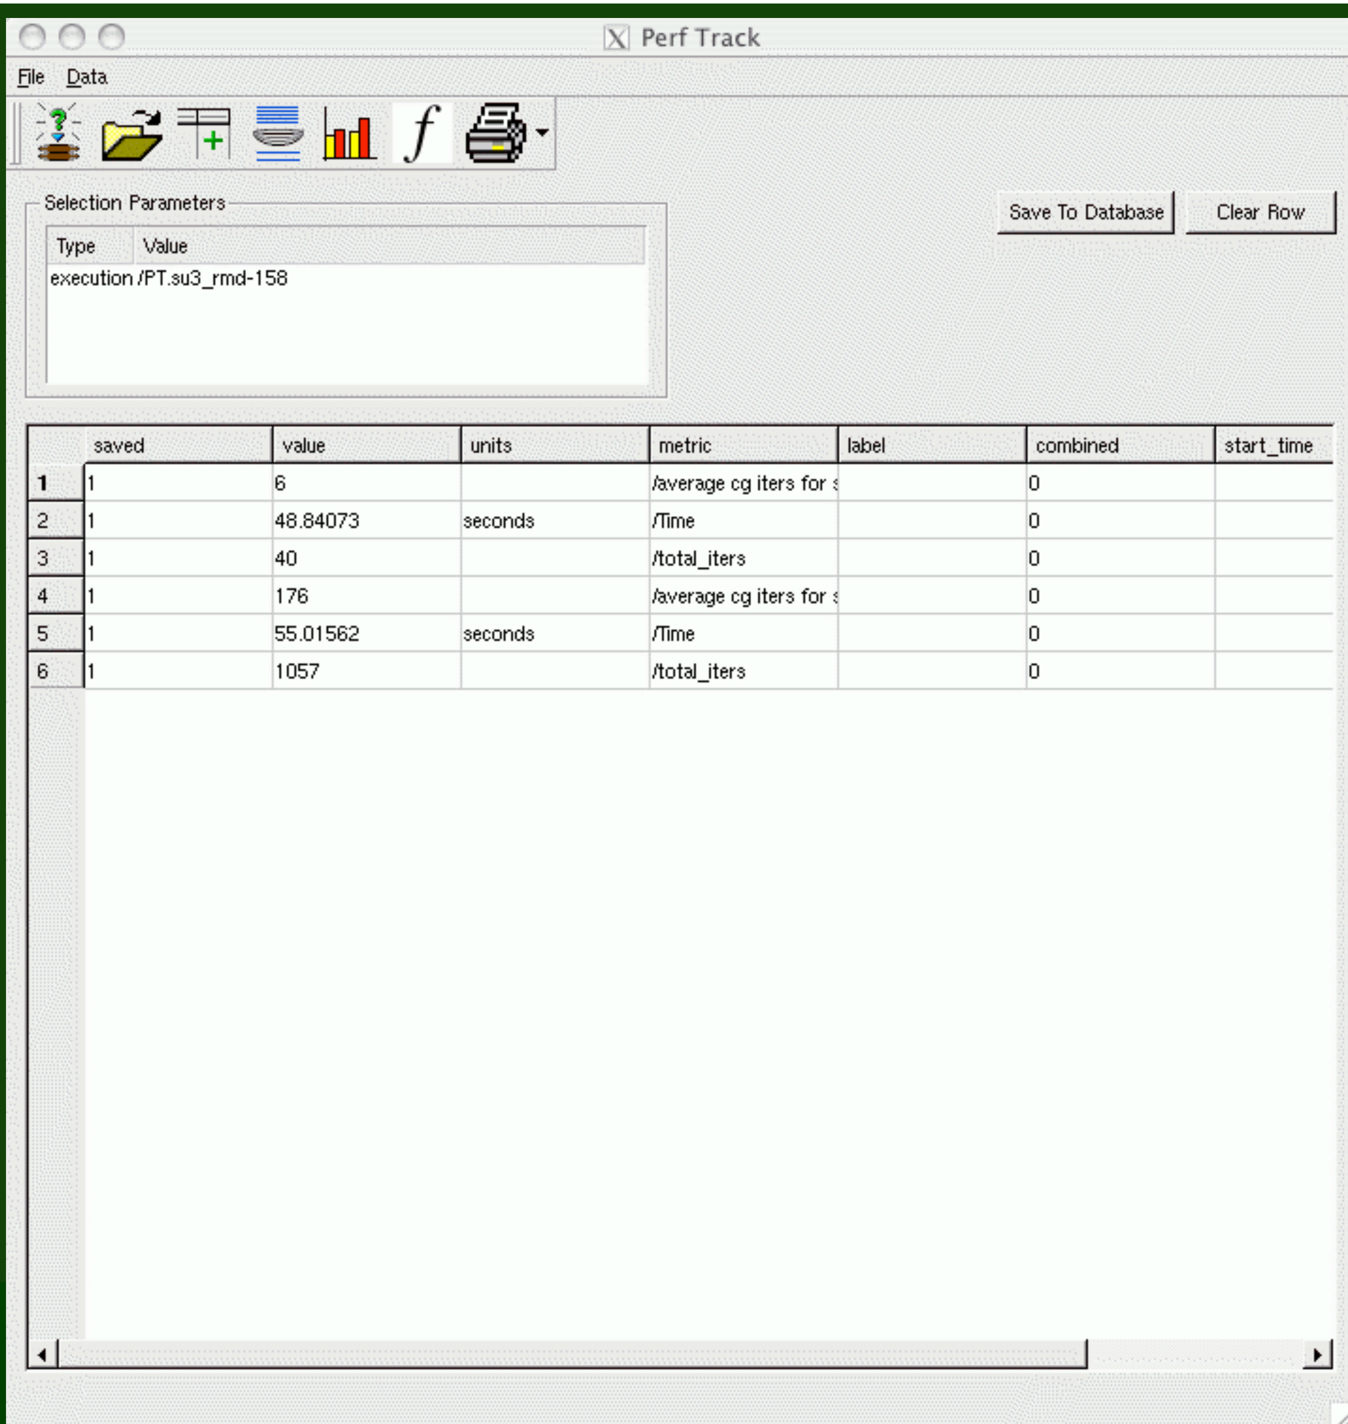

Type	Value
execution	/PT.su3_rmd-158

	saved	value	units	metric	label	combined	start_time
1	1	6		/average cg iters for s		0	
2	1	48.84073	seconds	/Time		0	
3	1	40		/total_iters		0	
4	1	176		/average cg iters for s		0	
5	1	55.01562	seconds	/Time		0	
6	1	1057		/total_iters		0	

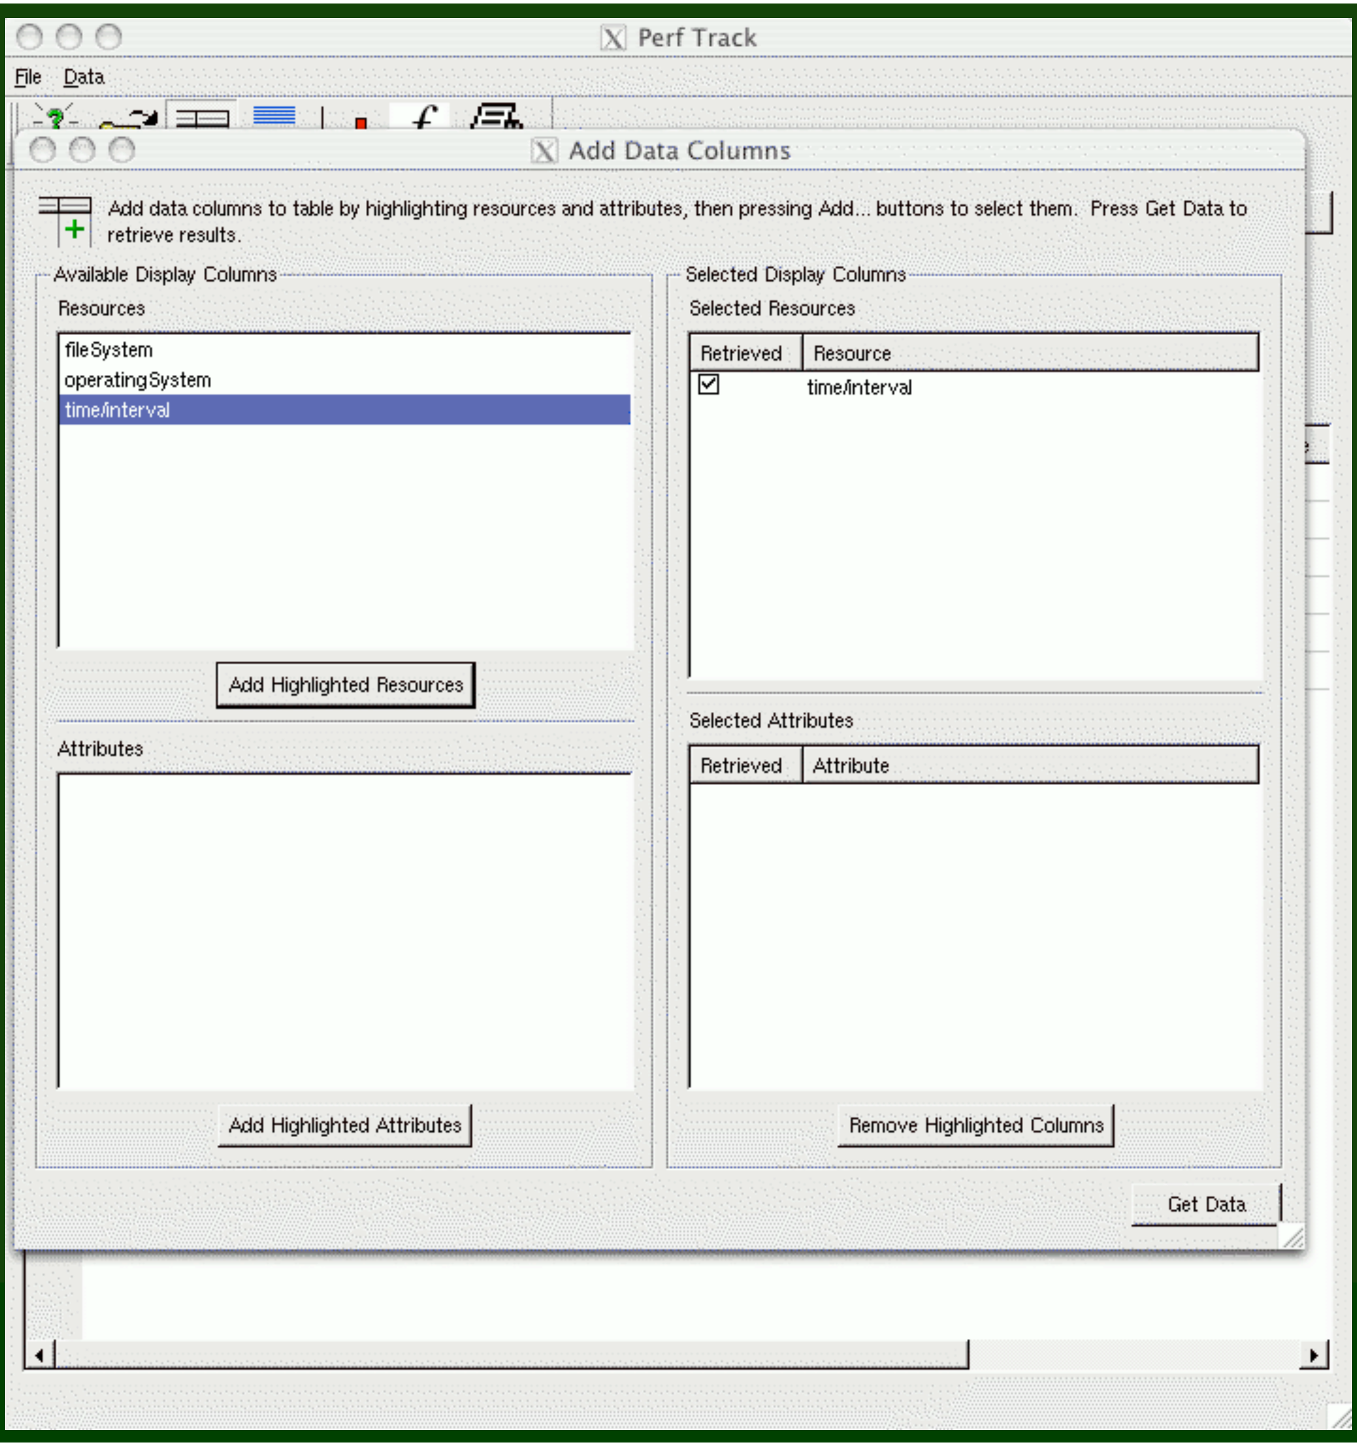

Perf Track
File
Data
Add Data Columns
Add data columns to table by highlighting resources and attributes, then pressing Add... buttons to select them. Press Get Data to retrieve results.
Available Display Columns
Resources
fileSystem
operatingSystem
time/interval
Add Highlighted Resources
Attributes
Add Highlighted Attributes
Selected Display Columns
Selected Resources
Retrieved
Resource
time/interval
Selected Attributes
Retrieved
Attribute
Remove Highlighted Columns
Get Data

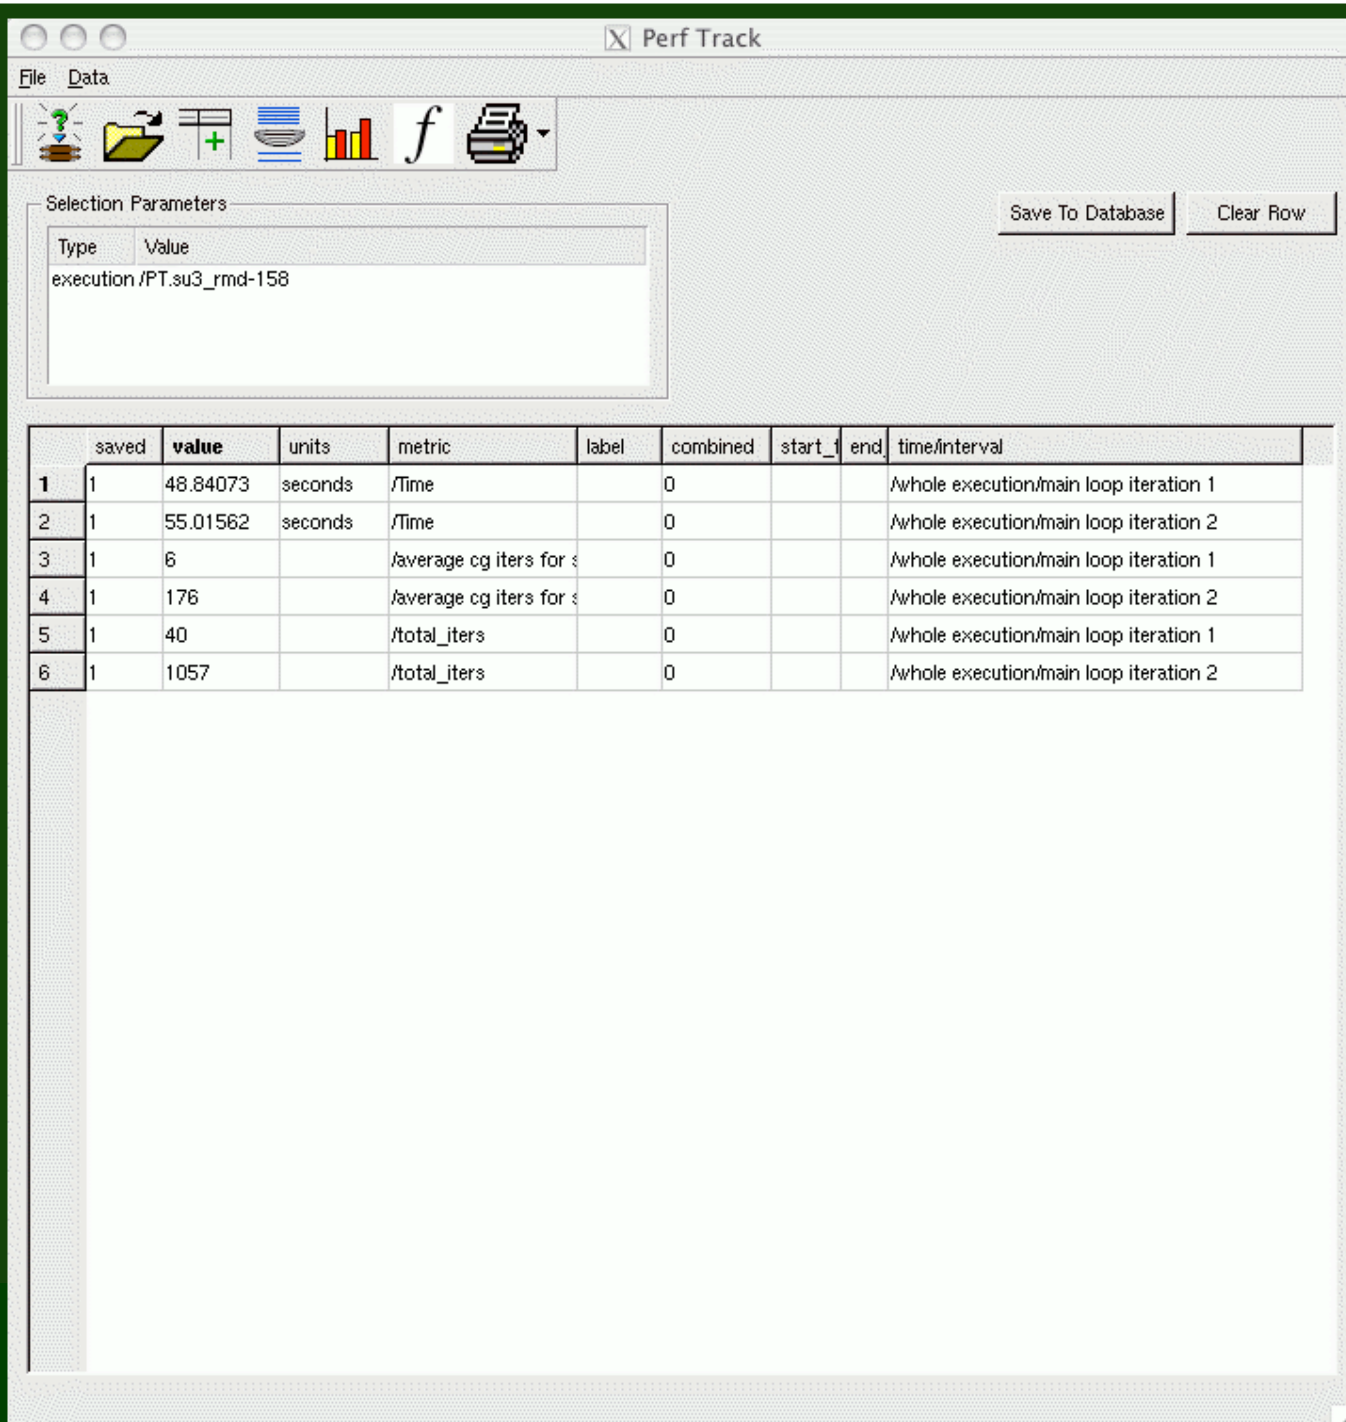
Perf Track
File Data
Selection Parameters
Type Value
execution /PT.su3_rmd-158
Save To Database
Clear Row
saved value units metric label combined start_t end time/interval
1 1 48.84073 seconds /Time 0 /whole execution/main loop iteration 1
2 1 55.01562 seconds /Time 0 /whole execution/main loop iteration 2
3 1 6 /average cg iters for s 0 /whole execution/main loop iteration 1
4 1 176 /average cg iters for s 0 /whole execution/main loop iteration 2
5 1 40 /total_iters 0 /whole execution/main loop iteration 1
6 1 1057 /total_iters 0 /whole execution/main loop iteration 2

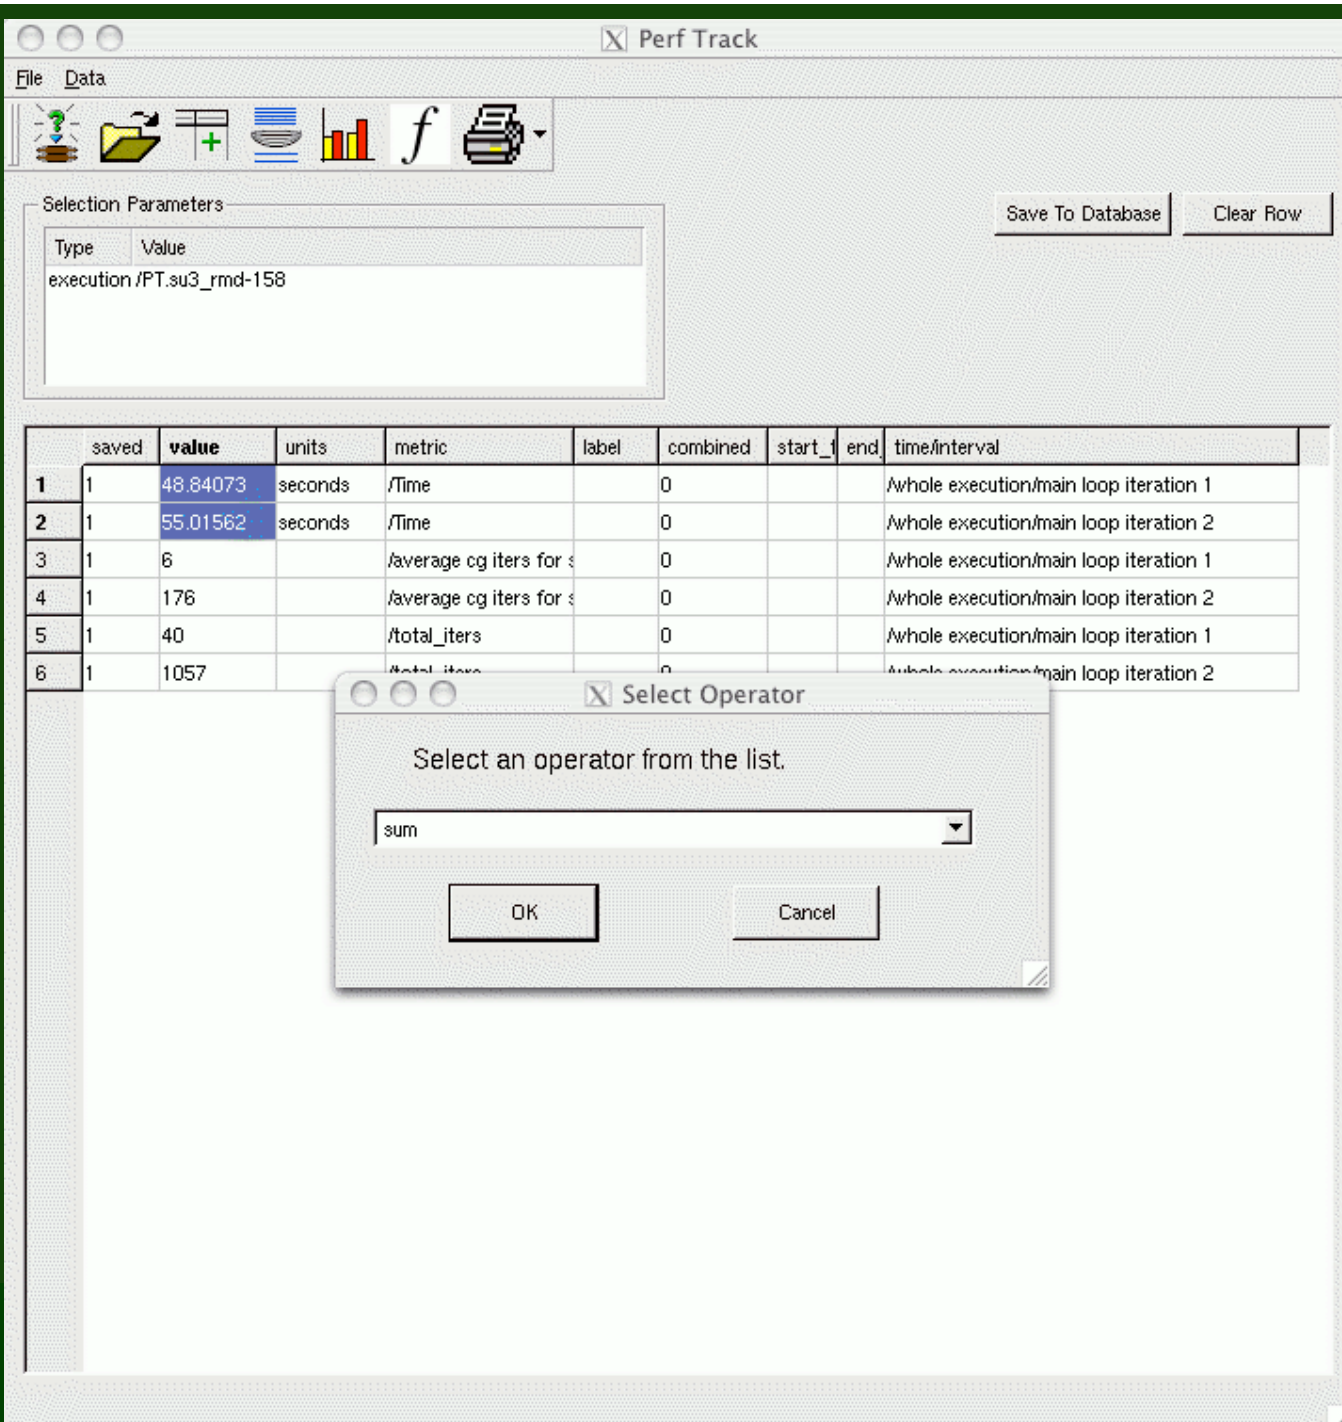
Perf Track
File Data
Selection Parameters
Type Value
execution /PT.su3_rmd-158
Save To Database
Clear Row
saved value units metric label combined start_t end time/interval
1 1 48.84073 seconds /Time 0 /whole execution/main loop iteration 1
2 1 55.01562 seconds /Time 0 /whole execution/main loop iteration 2
3 1 6 /average cg iters for s 0 /whole execution/main loop iteration 1
4 1 176 /average cg iters for s 0 /whole execution/main loop iteration 2
5 1 40 /total_iters 0 /whole execution/main loop iteration 1
6 1 1057 main loop iteration 2
Select Operator
Select an operator from the list.
sum
OK
Cancel

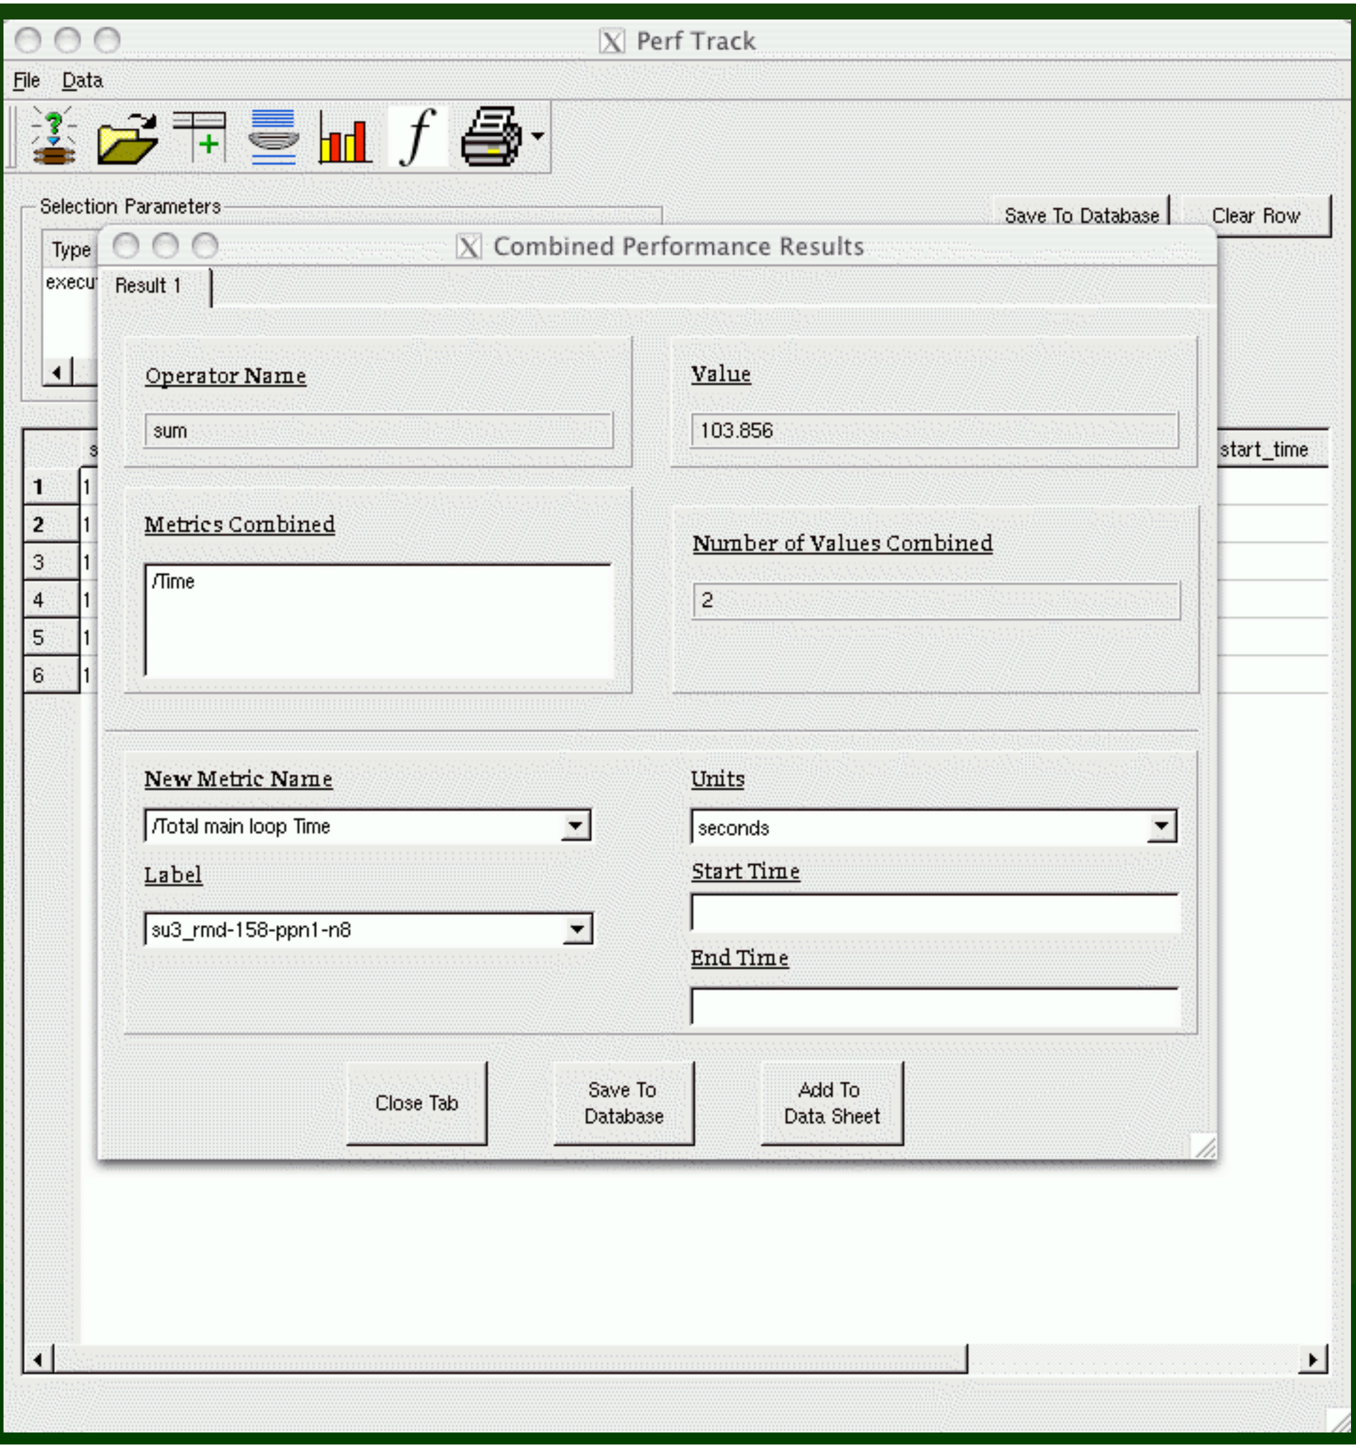
Perf Track
File
Data
Selection Parameters
Type
execu
Save To Database
Clear Row
start_time
Combined Performance Results
Result 1
Operator Name
sum
Value
103.856
Metrics Combined
/Time
Number of Values Combined
2
New Metric Name
/Total main loop Time
Units
seconds
Label
su3_rmd-158-ppn1-n8
Start Time
End Time
Close Tab
Save To Database
Add To Data Sheet

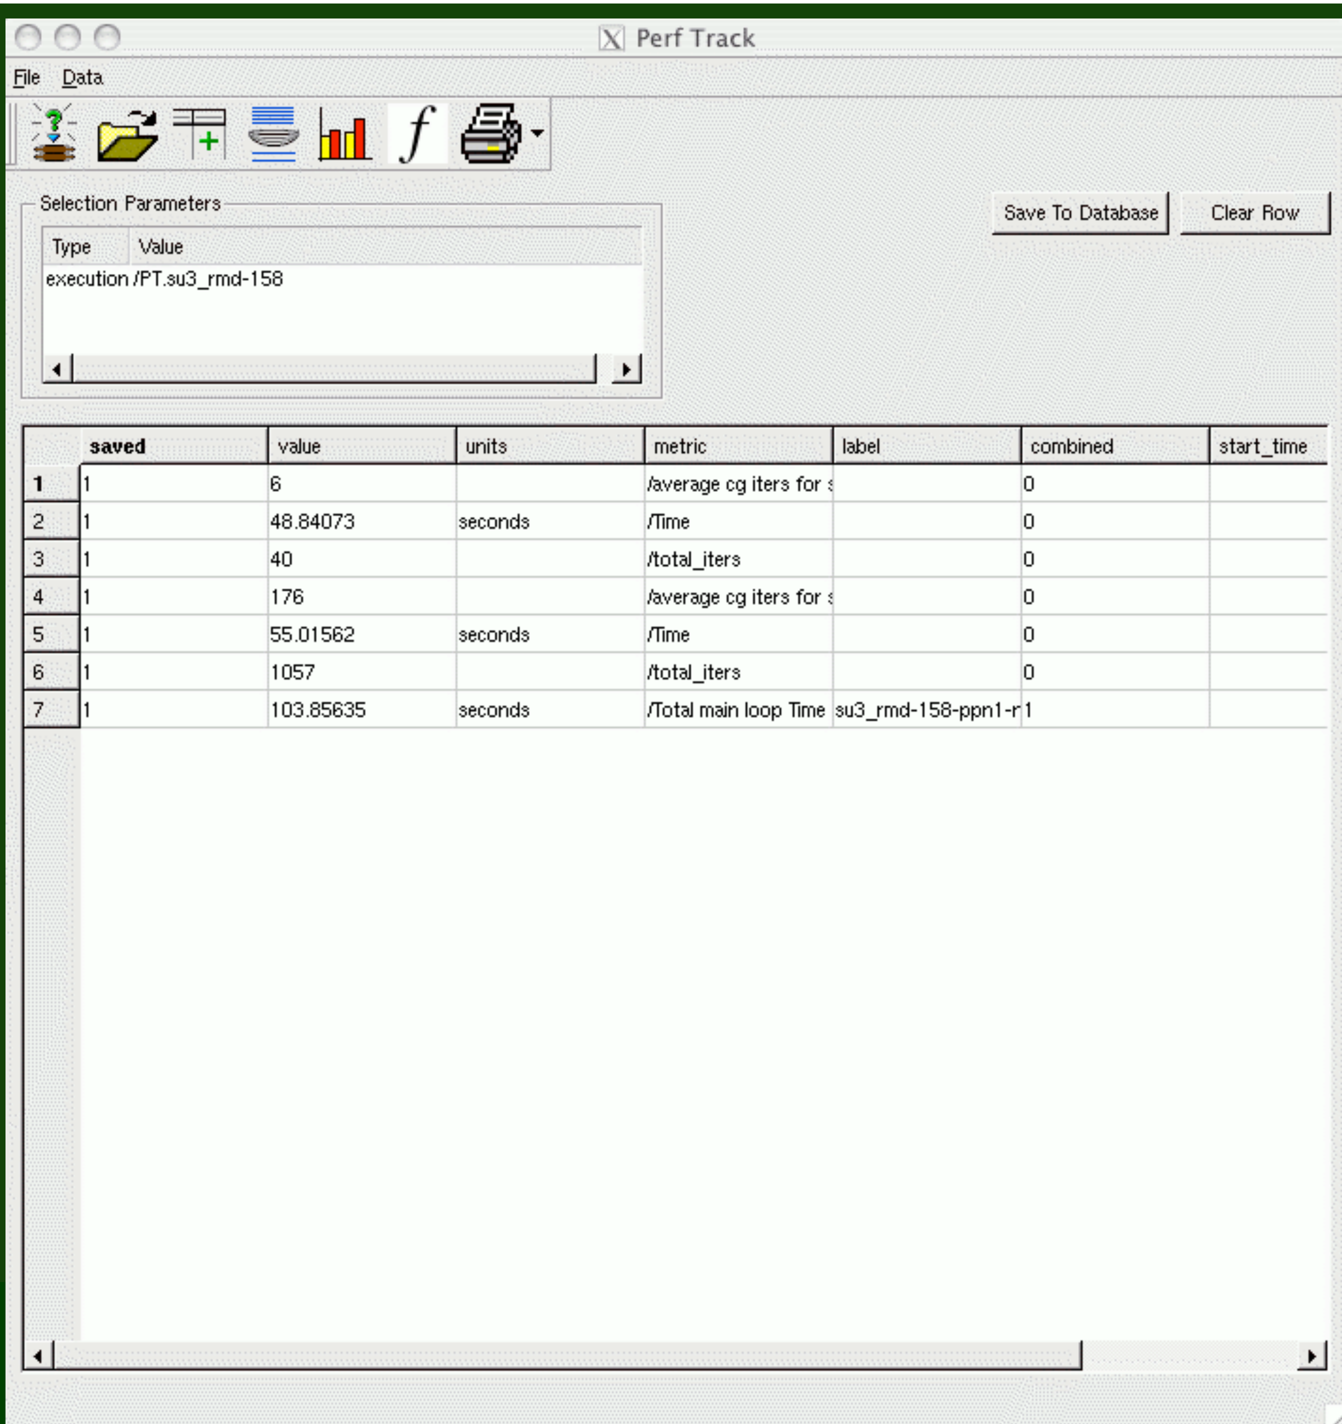
Perf Track
File
Data
Selection Parameters
Type
Value
execution /PT.su3_rmd-158
Save To Database
Clear Row
saved
value
units
metric
label
combined
start_time
1 1 6 /average cg iters for s 0
2 1 48.84073 seconds /Time 0
3 1 40 /total_iters 0
4 1 176 /average cg iters for s 0
5 1 55.01562 seconds /Time 0
6 1 1057 /total_iters 0
7 1 103.85635 seconds /Total main loop Time su3_rmd-158-ppn1-r 1

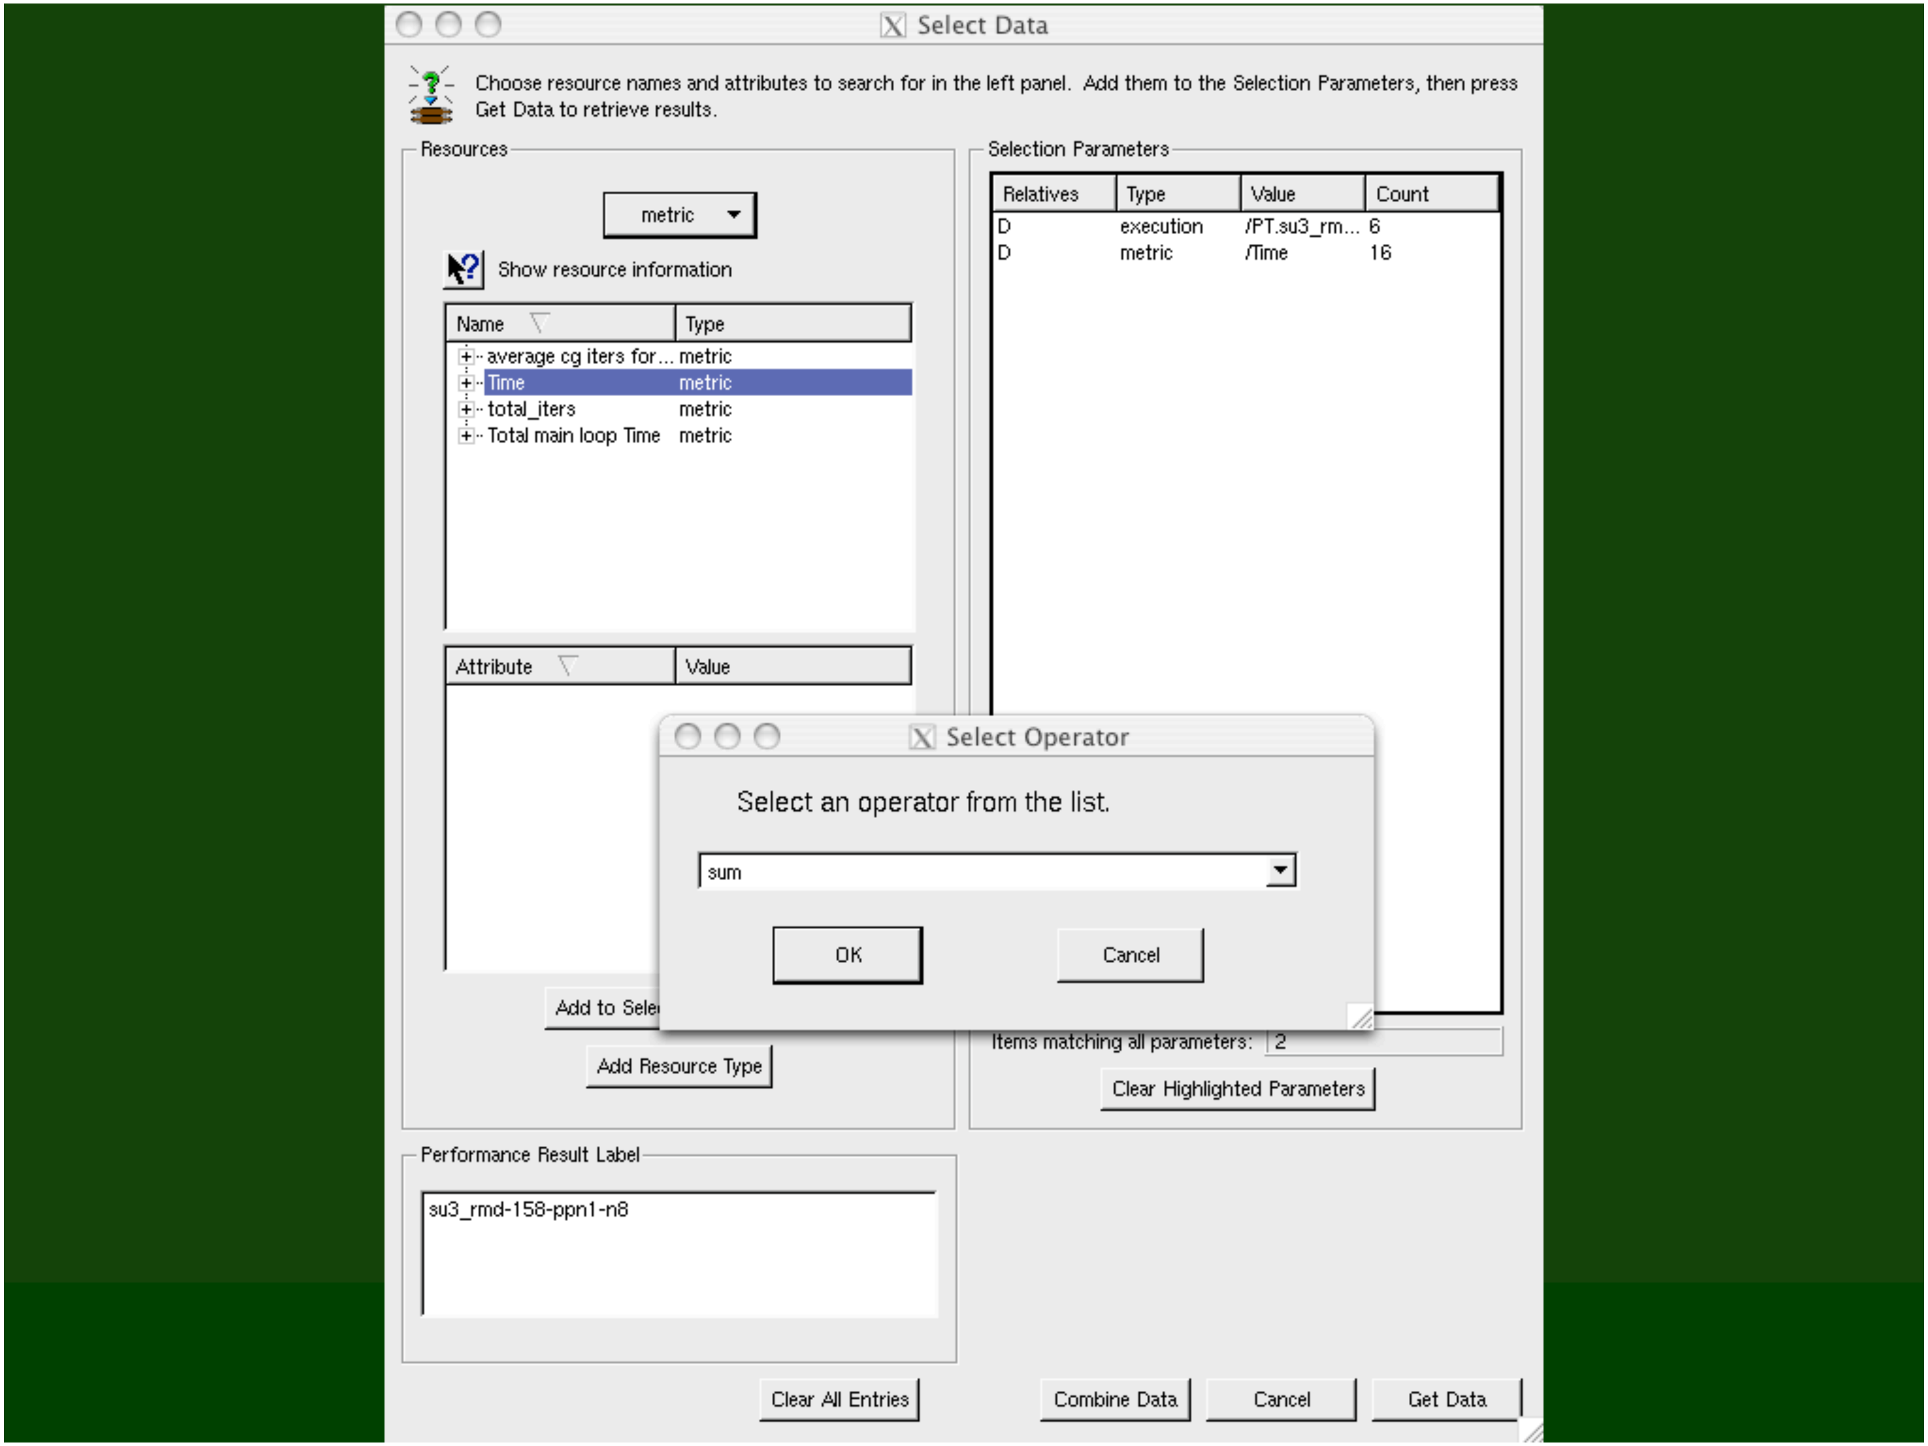
Select Data
Choose resource names and attributes to search for in the left panel. Add them to the Selection Parameters, then press Get Data to retrieve results.
Resources
metric
Show resource information
Name
Type
average cg iters for... metric
Time metric
total_iters metric
Total main loop Time metric
Attribute
Value
Add to Sele
Add Resource Type
Selection Parameters
Relatives
Type
Value
Count
D execution /PT.su3_rm... 6
D metric /Time 16
Items matching all parameters: 2
Clear Highlighted Parameters
Performance Result Label
su3_rmd-158-ppn1-n8
Clear All Entries
Combine Data
Cancel
Get Data
Select Operator
Select an operator from the list.
sum
OK
Cancel

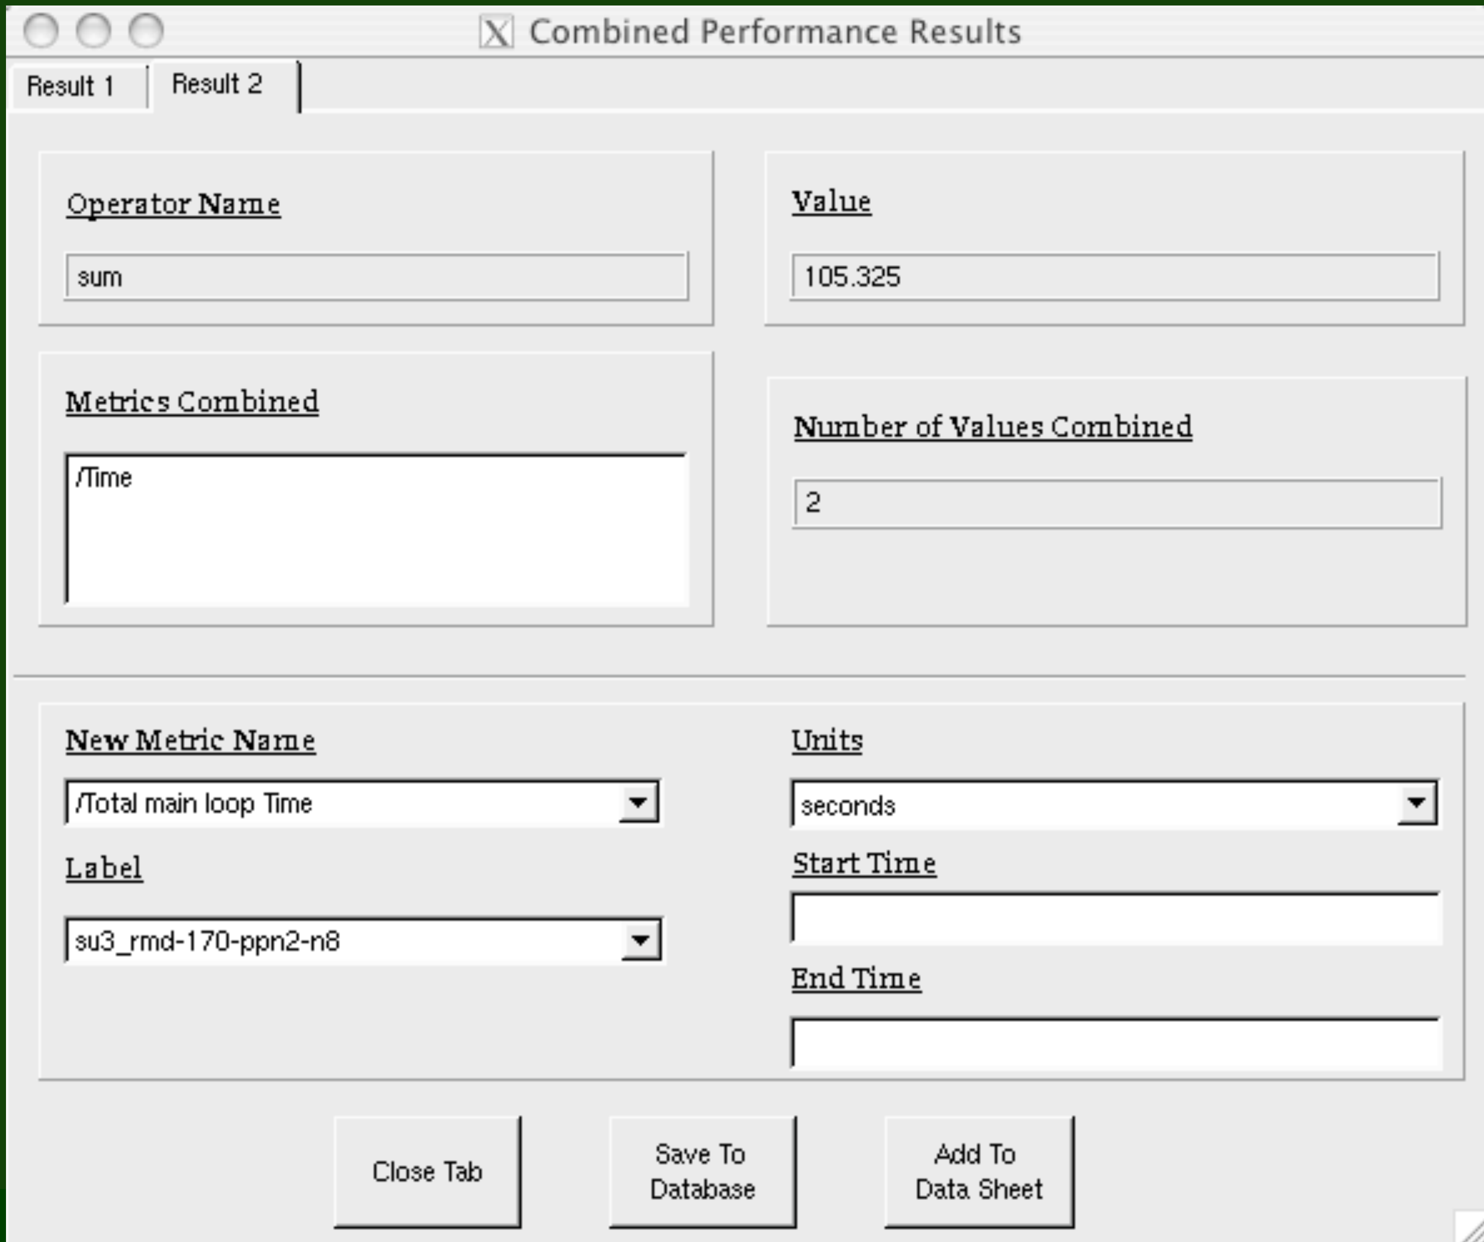
Combined Performance Results
Result 1
Result 2
Operator Name
sum
Value
105.325
Metrics Combined
/Time
Number of Values Combined
2
New Metric Name
/Total main loop Time
Units
seconds
Label
su3_rmd-170-ppn2-n8
Start Time
End Time
Close Tab
Save To Database
Add To Data Sheet

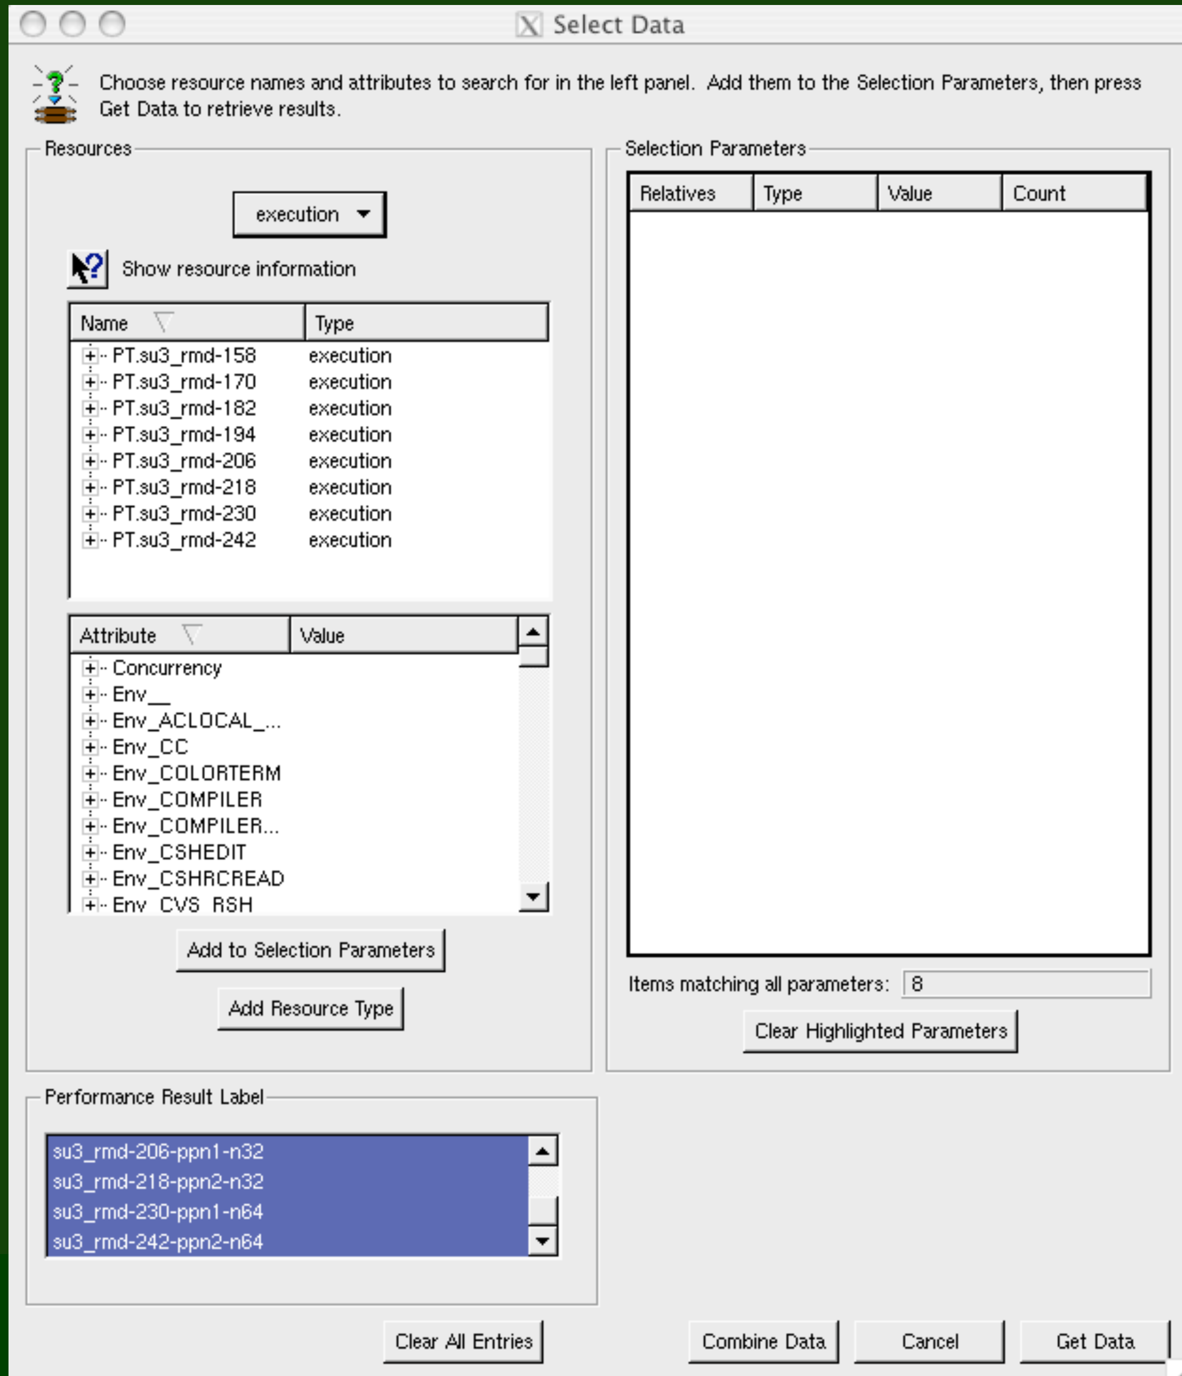

Select Data
Choose resource names and attributes to search for in the left panel. Add them to the Selection Parameters, then press Get Data to retrieve results.
Resources
execution
Show resource information
Name
Type
PT.su3_rmd-158 execution
PT.su3_rmd-170 execution
PT.su3_rmd-182 execution
PT.su3_rmd-194 execution
PT.su3_rmd-206 execution
PT.su3_rmd-218 execution
PT.su3_rmd-230 execution
PT.su3_rmd-242 execution
Attribute
Value
Concurrency
Env__
Env_ACLOCAL_...
Env_CC
Env_COLORTERM
Env_COMPILER
Env_COMPILER...
Env_CSHEDIT
Env_CSHRCREAD
Env_CVS_RSH
Add to Selection Parameters
Add Resource Type
Selection Parameters
Relatives
Type
Value
Count
Items matching all parameters: 8
Clear Highlighted Parameters
Performance Result Label
su3_rmd-206-ppn1-n32
su3_rmd-218-ppn2-n32
su3_rmd-230-ppn1-n64
su3_rmd-242-ppn2-n64
Clear All Entries
Combine Data
Cancel
Get Data

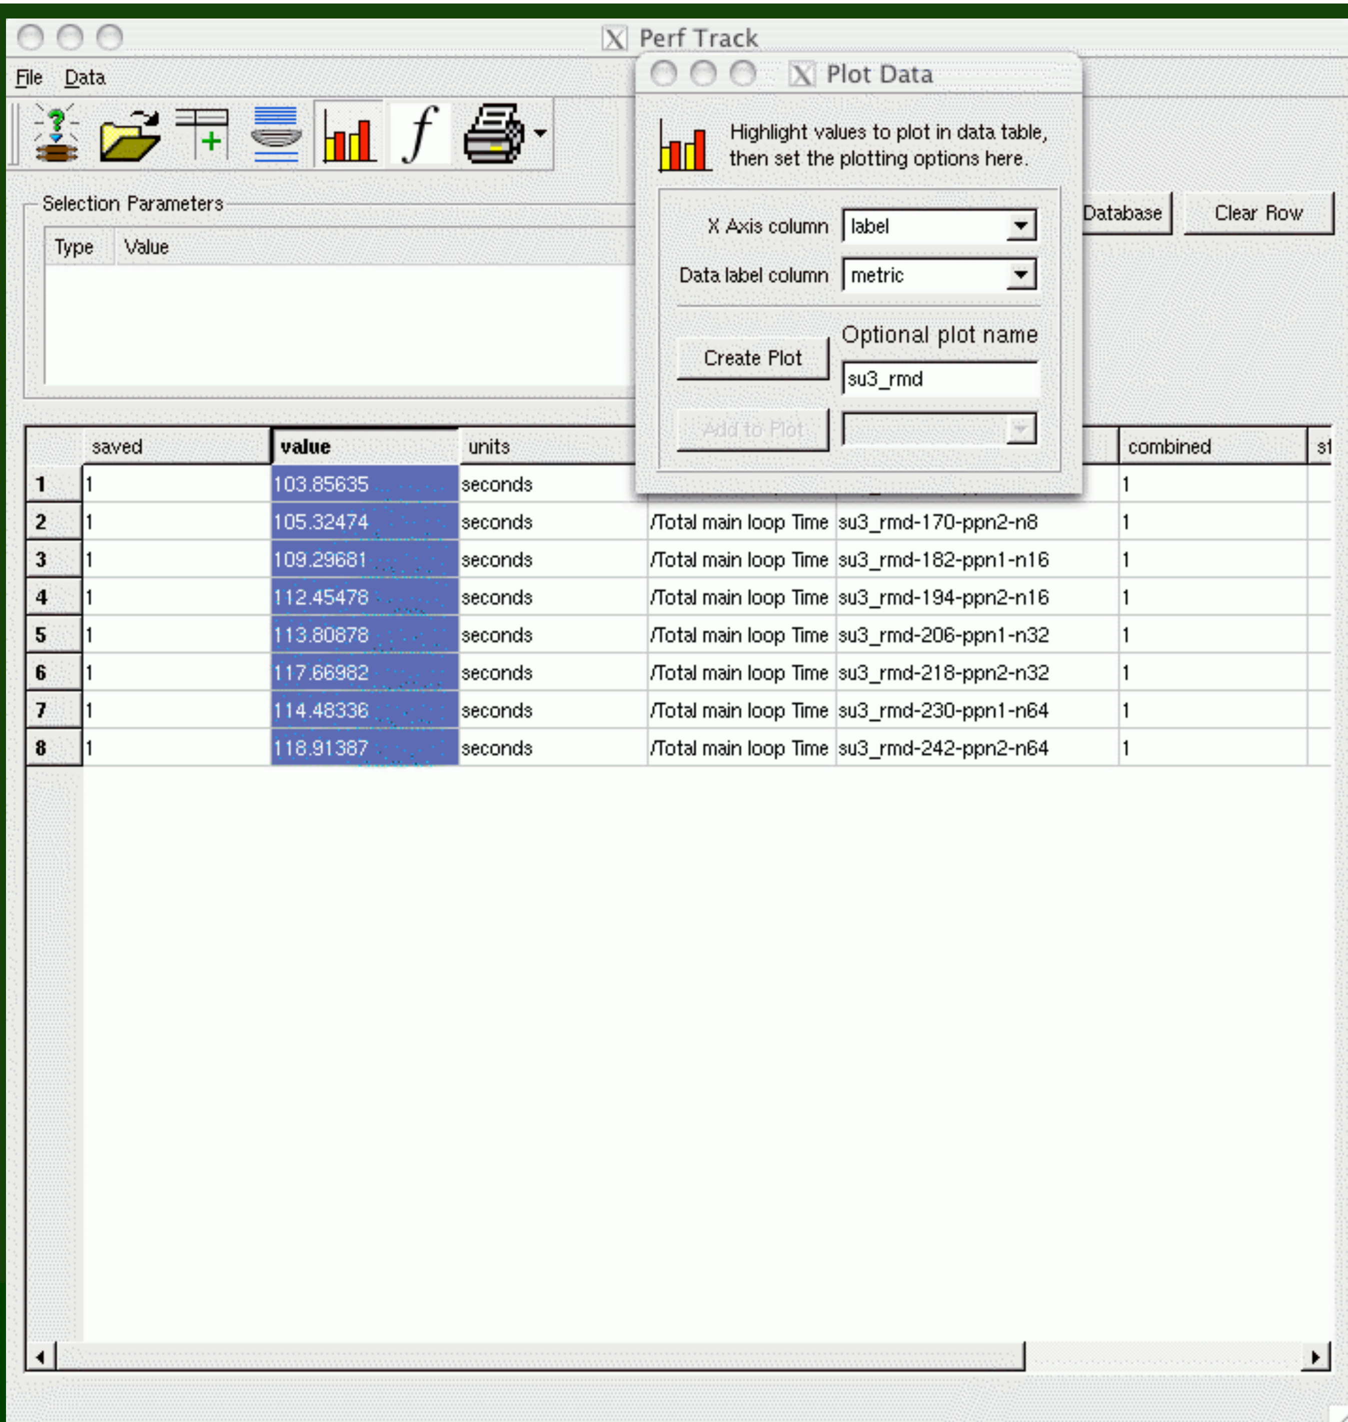
Perf Track
File Data
Plot Data
Highlight values to plot in data table, then set the plotting options here.
X Axis column label
Data label column metric
Create Plot
Optional plot name
su3_rmd
Add to Plot
Database
Clear Row
Selection Parameters
Type Value
saved
value
units
combined
1 1 103.85635 seconds 1
2 1 105.32474 seconds /Total main loop Time su3_rmd-170-ppn2-n8 1
3 1 109.29681 seconds /Total main loop Time su3_rmd-182-ppn1-n16 1
4 1 112.45478 seconds /Total main loop Time su3_rmd-194-ppn2-n16 1
5 1 113.80878 seconds /Total main loop Time su3_rmd-206-ppn1-n32 1
6 1 117.66982 seconds /Total main loop Time su3_rmd-218-ppn2-n32 1
7 1 114.48336 seconds /Total main loop Time su3_rmd-230-ppn1-n64 1
8 1 118.91387 seconds /Total main loop Time su3_rmd-242-ppn2-n64 1

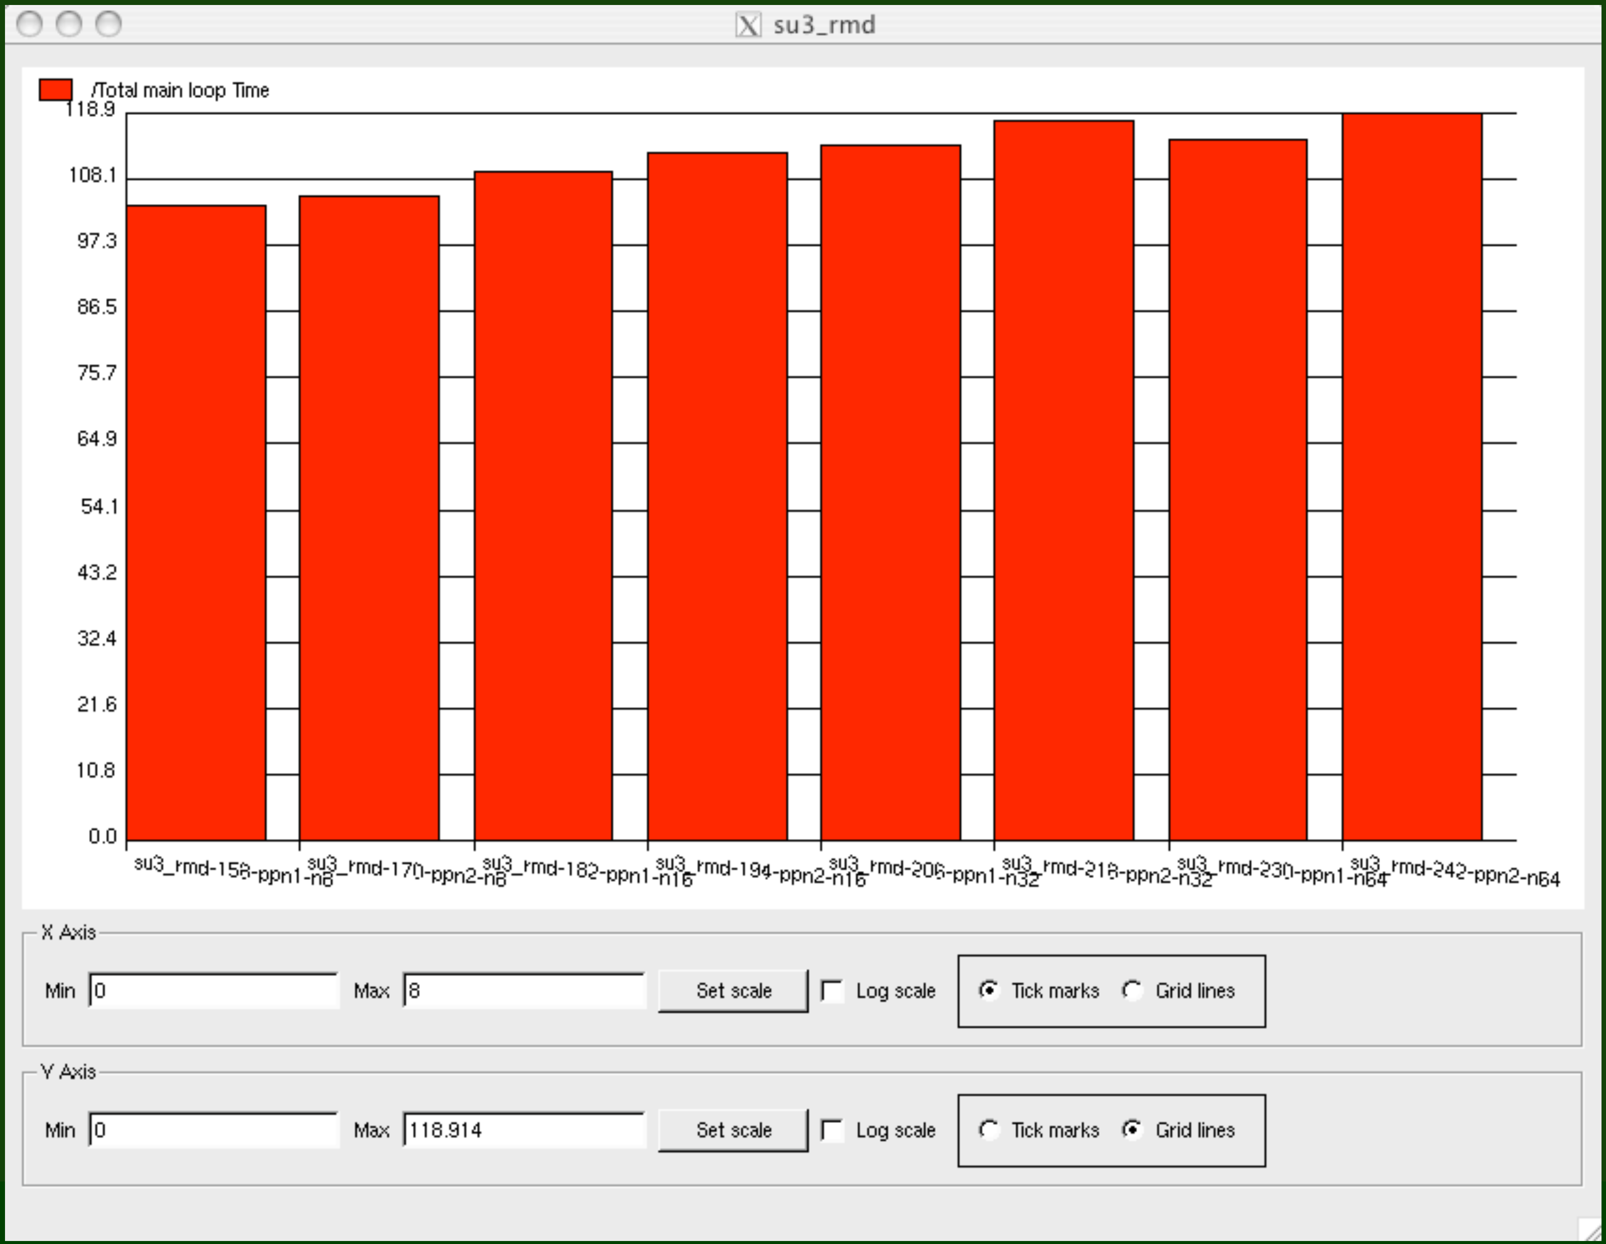
su3_rmd
/Total main loop Time
118.9
108.1
97.3
86.5
75.7
64.9
54.1
43.2
32.4
21.6
10.8
0.0
su3_rmd-158-ppn1-n8
su3_rmd-170-ppn2-n8
su3_rmd-182-ppn1-n16
su3_rmd-194-ppn2-n16
su3_rmd-206-ppn1-n32
su3_rmd-218-ppn2-n32
su3_rmd-230-ppn1-n64
su3_rmd-242-ppn2-n64
X Axis
Min 0
Max 8
Set scale
Log scale
Tick marks
Grid lines
Y Axis
Min 0
Max 118.914
Set scale
Log scale
Tick marks
Grid lines

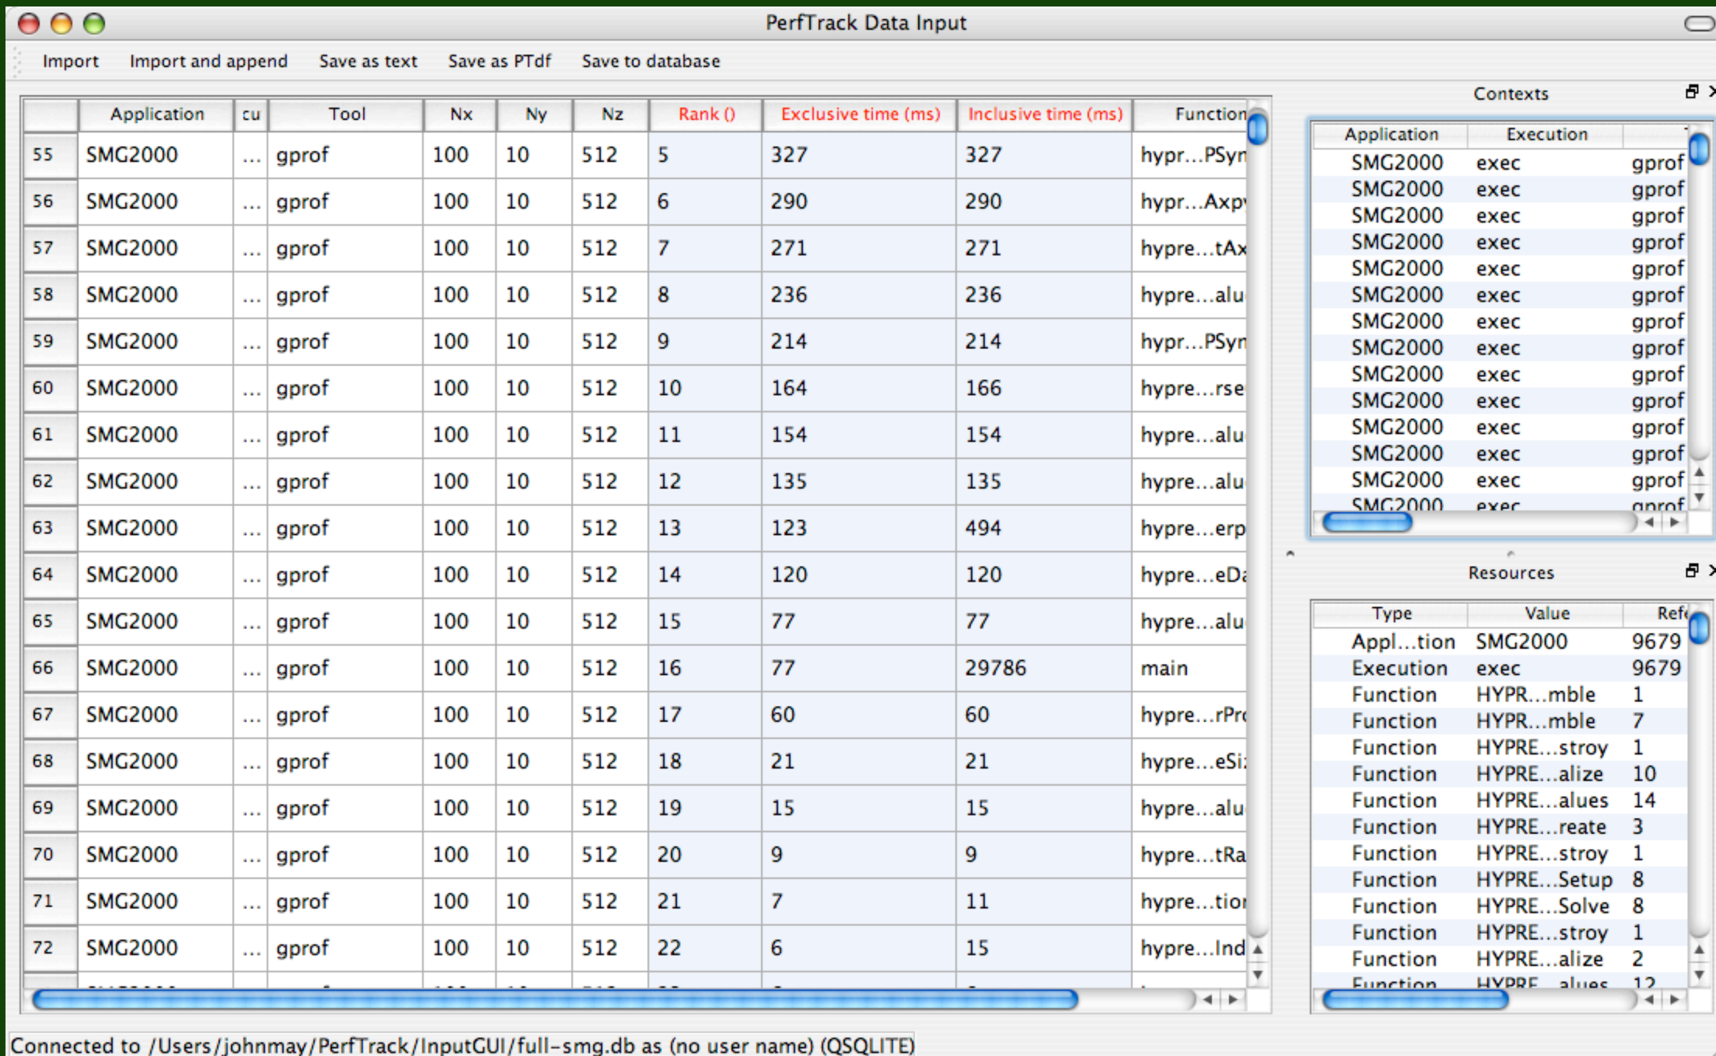

	Application	cu	Tool	Nx	Ny	Nz	Rank ()	Exclusive time (ms)	Inclusive time (ms)	Function
55	SMG2000	...	gprof	100	10	512	5	327	327	hypr...PSyn
56	SMG2000	...	gprof	100	10	512	6	290	290	hypr...Axpy
57	SMG2000	...	gprof	100	10	512	7	271	271	hypre...tAx
58	SMG2000	...	gprof	100	10	512	8	236	236	hypre...alu
59	SMG2000	...	gprof	100	10	512	9	214	214	hypr...PSyn
60	SMG2000	...	gprof	100	10	512	10	164	166	hypre...rse
61	SMG2000	...	gprof	100	10	512	11	154	154	hypre...alu
62	SMG2000	...	gprof	100	10	512	12	135	135	hypre...alu
63	SMG2000	...	gprof	100	10	512	13	123	494	hypre...erp
64	SMG2000	...	gprof	100	10	512	14	120	120	hypre...eDa
65	SMG2000	...	gprof	100	10	512	15	77	77	hypre...alu
66	SMG2000	...	gprof	100	10	512	16	77	29786	main
67	SMG2000	...	gprof	100	10	512	17	60	60	hypre...rPro
68	SMG2000	...	gprof	100	10	512	18	21	21	hypre...eSiz
69	SMG2000	...	gprof	100	10	512	19	15	15	hypre...alu
70	SMG2000	...	gprof	100	10	512	20	9	9	hypre...tRa
71	SMG2000	...	gprof	100	10	512	21	7	11	hypre...tion
72	SMG2000	...	gprof	100	10	512	22	6	15	hypre...Ind

Application	Execution	
SMG2000	exec	gprof
SMG2000	exec	gprof
SMG2000	exec	gprof
SMG2000	exec	gprof
SMG2000	exec	gprof
SMG2000	exec	gprof
SMG2000	exec	gprof
SMG2000	exec	gprof
SMG2000	exec	gprof
SMG2000	exec	gprof
SMG2000	exec	gprof
SMG2000	exec	gprof
SMG2000	exec	gprof

Type	Value	Ref
Appl...tion	SMG2000	9679
Execution	exec	9679
Function	HYPR...mble	1
Function	HYPR...mble	7
Function	HYPRE...stroy	1
Function	HYPRE...alize	10
Function	HYPRE...alues	14
Function	HYPRE...reate	3
Function	HYPRE...stroy	1
Function	HYPRE...Setup	8
Function	HYPRE...Solve	8
Function	HYPRE...stroy	1
Function	HYPRE...alize	2

PerfTrack

"Metadata Collection for Performance Analysis"
Karen L. Karavanic July 17, 2007

Portland State UNIVERSITY

How do people want to use PerfTrack?

- Traditional Performance Analysis/Tuning
- Fully automated performance regression testing
- Comparative evaluation of new platforms vs. old
- Effects of hardware and software upgrades
- OS kernel performance study
- Organized store to replace scattered files
- Sharing a single data store in collaborative studies

What do people want to store in PerfTrack ?

- All Performance Experiment Artifacts
  - HW counter data, profile data, trace data, benchmark output
  - "Barry's World": create the graph, save the graph, save the steps to create the graph
  - Paradyn artifacts: Search History Graph, Call graph, performance data histograms
  - Data from all common tools -- OpenSpeedshop, TAU, etc.
  - As much description of the build and runtime environments as possible

Machine Data Collection

- "Automated System Environment Capture For PerfTrack" Capstone Project Team: Aaron Amauba, Dave Vu, Steve Wooster
- What to collect?
- When to collect?
- We already have this information… right??
- Who knows?
- Device model number example
- Is the execution timestamp enough?

Machine Data Collection: Host System method

- Idea
  - Write scripts to run on the host system and directly measure the environment
- Features
  - Modular structure. User can specify tailor made modules for the system they are interested in.
  - Can be run anytime. Just kick off the script.
  - Allow the user to provide resource hierarchy information when a connection to the database is not available.
- Status
  - Currently implemented and tested for Linux

Machine Description -- Scaling Challenges

- Automated Host Data Collection eliminates tedious manual entry

BUT

- When to scan? BG/L -- 100k+ nodes; update frequency ??
  - Per Execution?? Weekly? Daily? Hourly?
- Who scans?
  - Each researcher ?
    - (requires running a "scan" program on each node, can be difficult to get these types of jobs scheduled by scheduler)
    - Do we really need all those copies??
- Delay in entry -> overwrite new data with old?? names??
- The best answer will involve lab support!

Collaborative Data Stores

- PERI-DB (Shirley Moore)
- Goal: Develop and deploy a data store for performance data sets of PERI project researchers
- Approach
  - Define PERI XML schema
  - Tools provide a mapping to/from PERI XML
- We extended our data collection scripts to output PERI xml
- In progress: translating PTdf <=> PERI xml
  - conversion of PTdf files to PERI xml
  - conversion of PERI xml files to PTdf
  - input of PERI xml to PerfTrack database
  - export of PERI xml from queries to PerfTrack database

Key Issues in Metadata Collection

- Rich Data Sets
  - build, platform, runtime environment, performance data
  - sparse data will impact results -- eg clustering
- Scalability
  - collection frequency: each run? each experiment? each user? each boot? each upgrade?
  - attributes for groups of resources
  - Time: when is new knowledge created?
  - new machine resource every year
  - attributes of machine resource change every few weeks
  - input data changes every few runs
  - runtime environment may change during one execution
  - diagnoses and comparative data during analysis

Key Issues in Metadata Collection

- Collaborative Data Stores
  - Need to map resources and results between local sites
- Porting difficulties
  - Commands for collecting metadata information vary from platform to platform
- Lack of common interfaces
  - e.g. file system software version information on Linux: Lustre vs GPFS

The PerfTrack Project

- http://www.llnl.gov/casc/perfTrack/
- http://www.cs.pdx.edu/~karavan/perftrack
- Karen Karavanic: karavan@cs.pdx.edu John May: johnmay@llnl.gov
- Karen L. Karavanic, John May, et al, "Integrating Database Technology with Comparison-based Parallel Performance Diagnosis: The PerfTrack Performance Experiment Management Tool," SC2005, November 2005, Seattle, WA.
- Kathryn Mohror, Rashawn Knapp, Nagalaxmi Karumbunathan
- This research supported in part by UC/LLNL subcontract #B539302.
- Portions of this work were performed under the auspices of the U.S. Department of Energy by the University of California Lawrence Livermore National Laboratory under contract No. W-7405-Eng-48.
- Clustering interfaces: Thomas Conerly, Abraham Neben (Portland Saturday Academy Internship Program for high school students)
- PPerfGrid: John Hoffman (PSU masters thesis)
- Capstone Project: Aaron Amauba, Dave Vu, Steve Wooster (PSU undergraduate course)