Testing Proof on Xen

CERN openlab

for DataGrid applications

Rosa M Garcia
Sverre Jarp

Index

- Proof communication
- Different scenarios
- Results on PC with processes
- Results on PC with domains (Xen)
- Comparison
- Conclusions

Proof communication

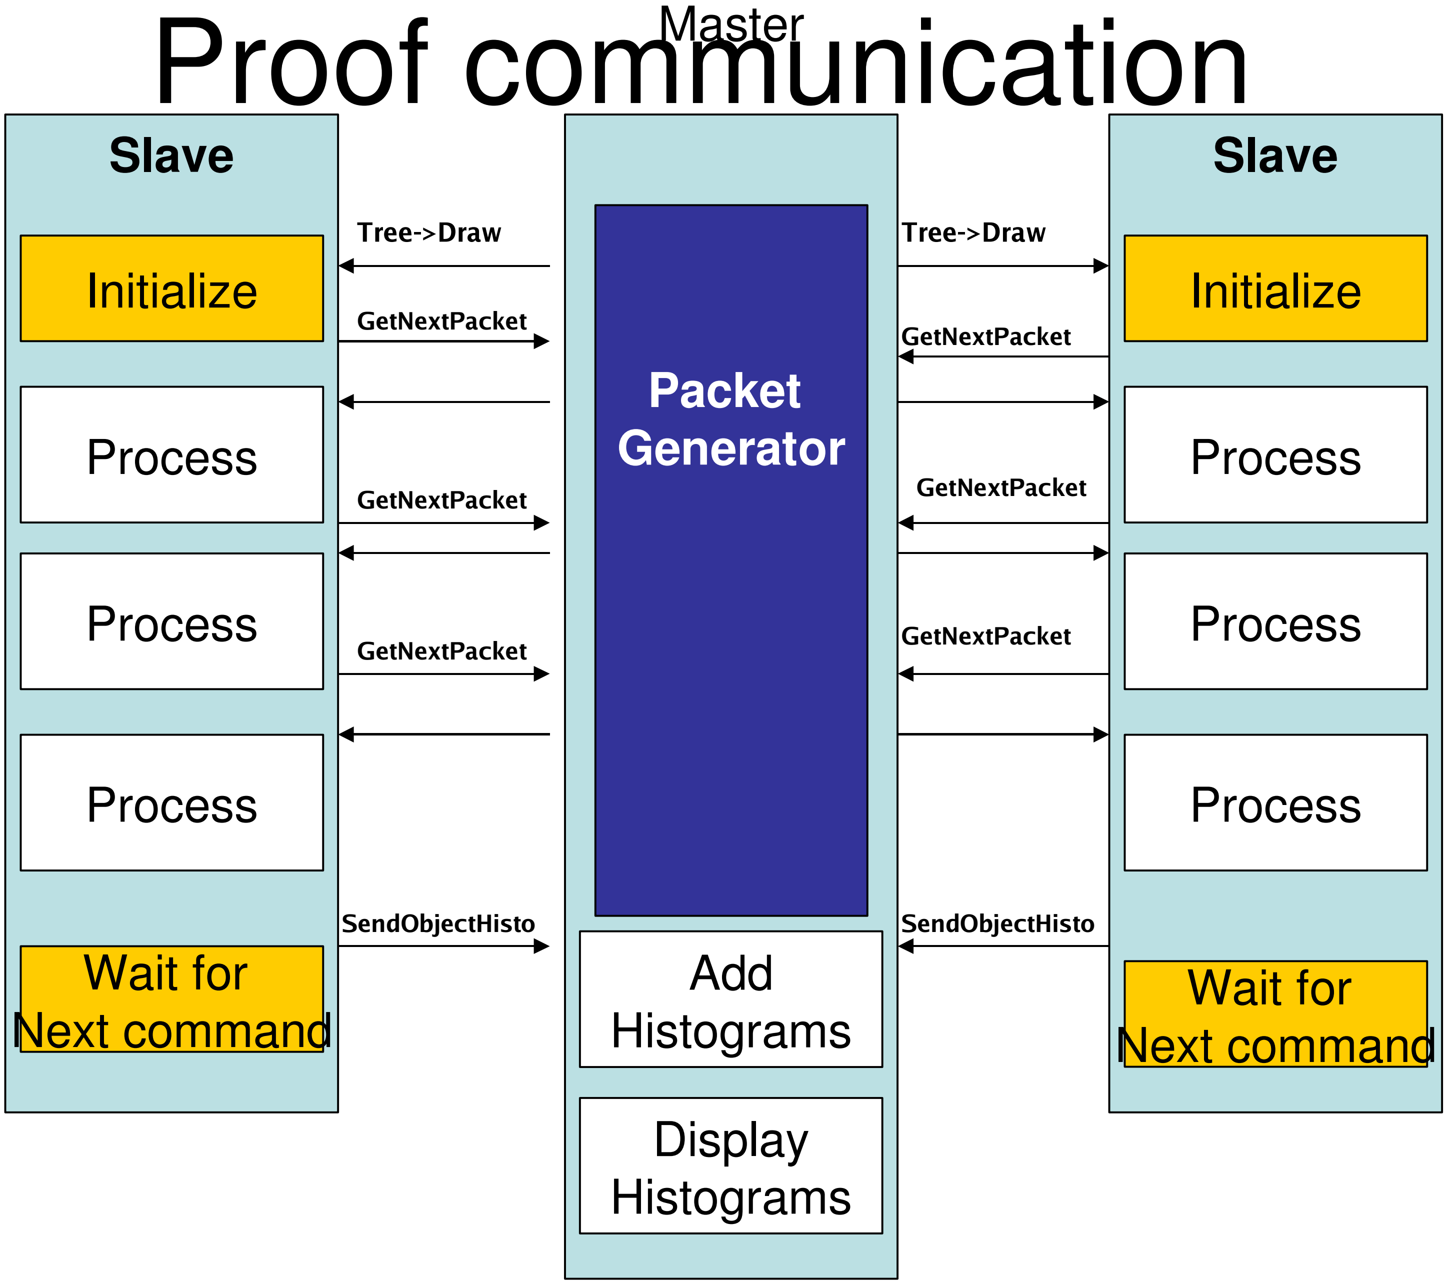

Different test scenarios

- 4 slaves + master distributed in 2 Pc running as processes

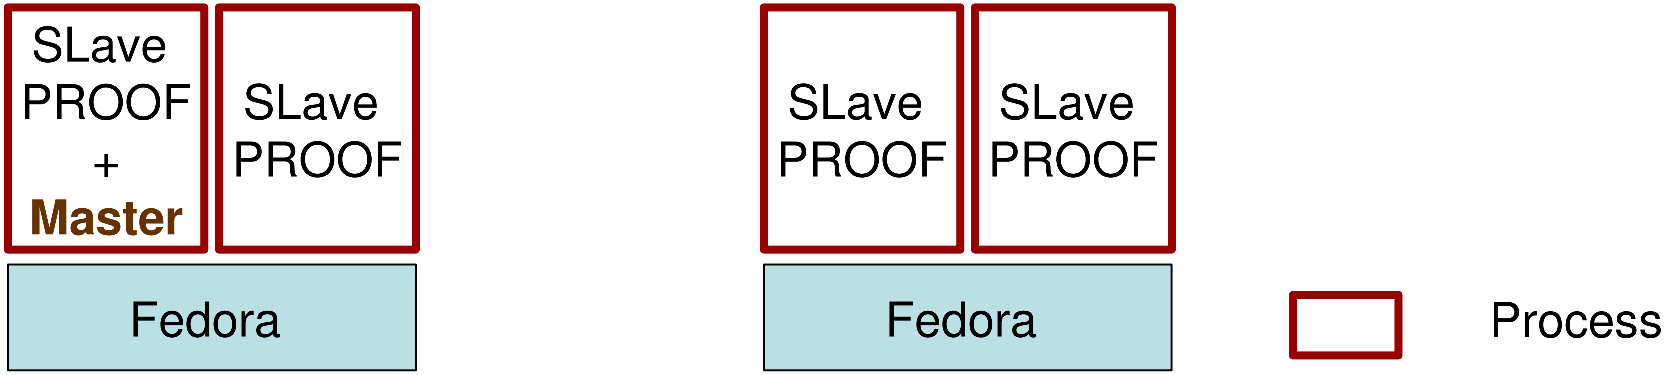

SLave PROOF + **Master**

SLave PROOF

Fedora

SLave PROOF

SLave PROOF

Fedora

Process

- 4 slaves + master running in independent domains

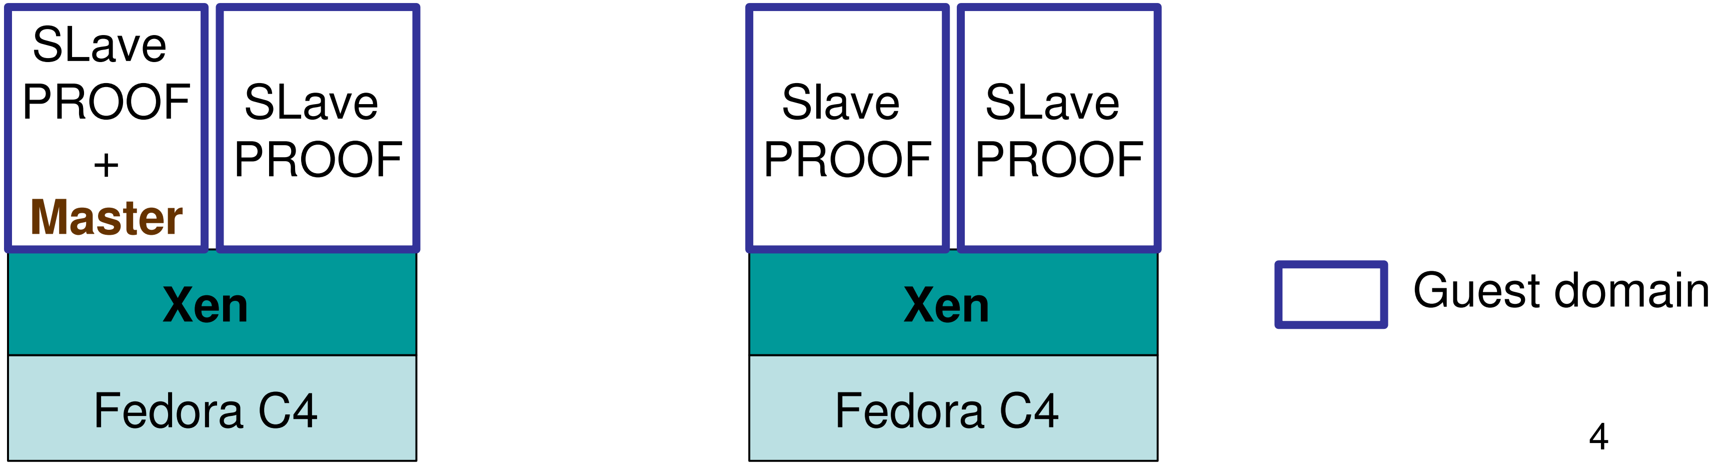

SLave PROOF + **Master**

SLave PROOF

Xen

Fedora C4

Slave PROOF

SLave PROOF

Xen

Fedora C4

Guest domain

Benchmark

- The benchmark provides 2 selectors, each one reading a different amount of data
  - ProcOpt → Reads 25% of the data.
  - Proc → Reads all the data
- Measurements in the Slave
  - Packet processing time
  - Get Packet Latency
  - CPU Time

First results in scenario 1

Benchmark: Reads only 25% of the data

Scenario: 2 PC with each slave is running in a process

The graphics include error bars per value

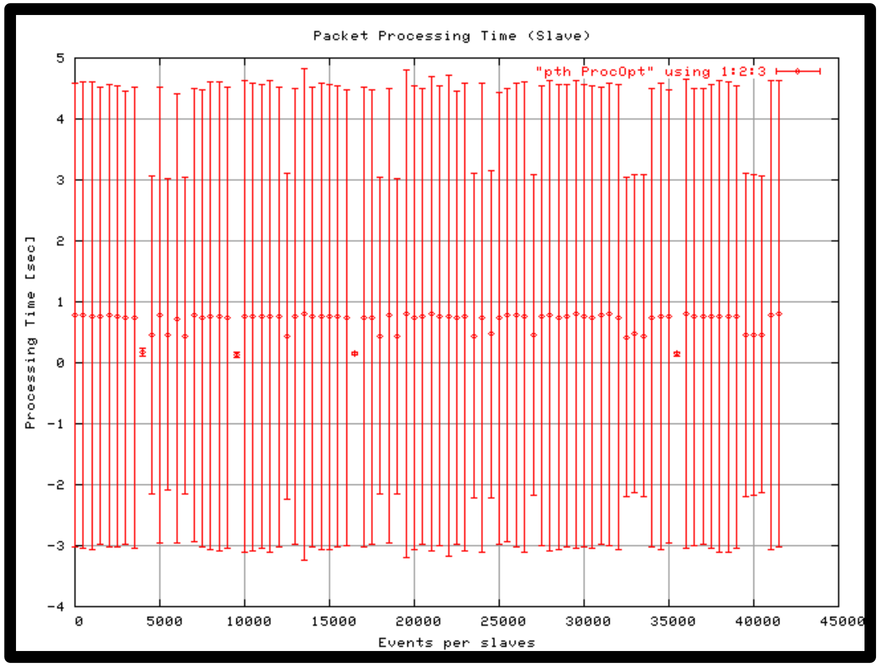

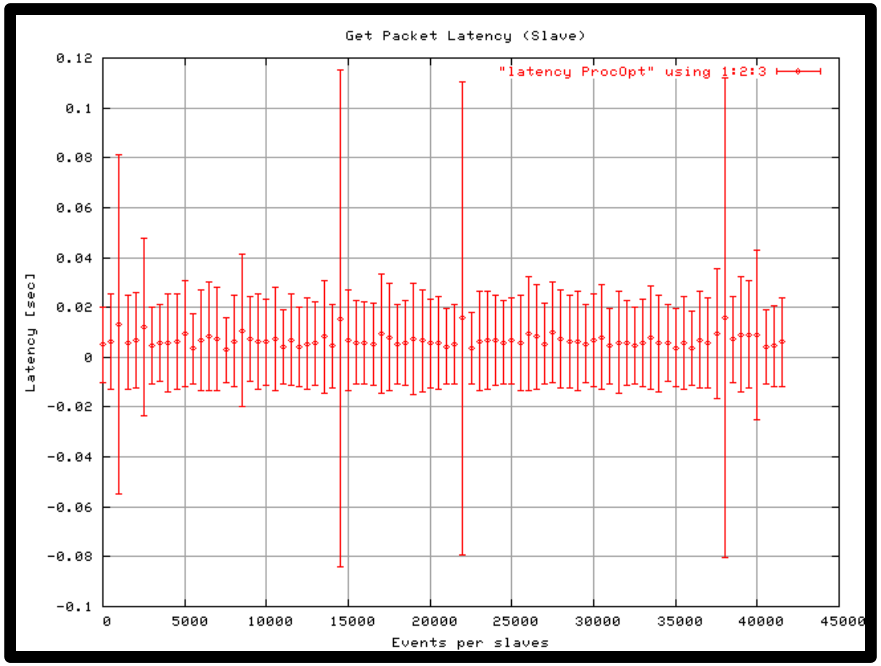

First results in scenario 1 (CPU Time)

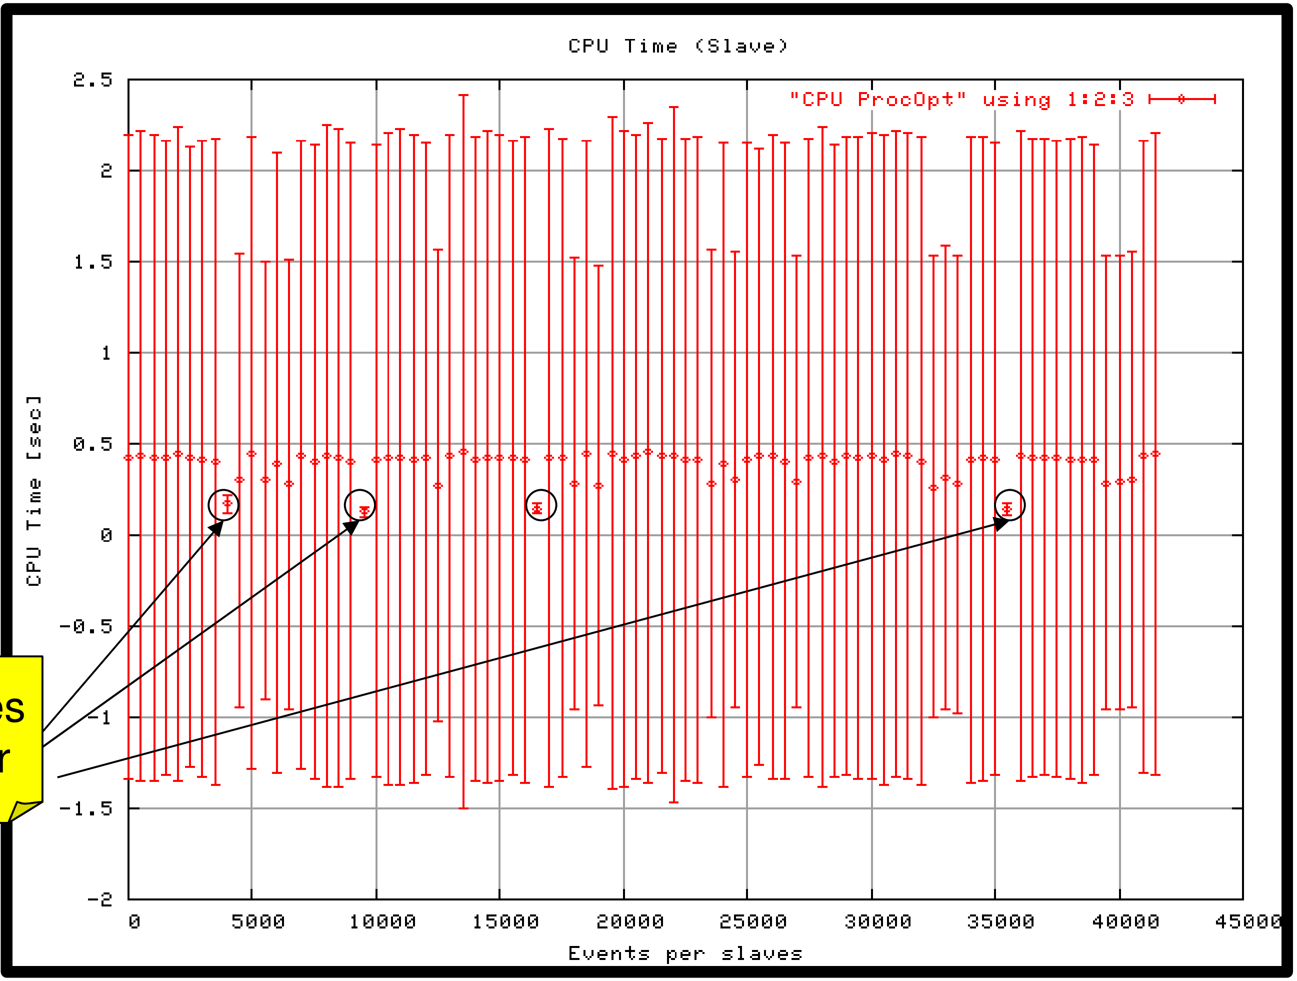

Good values
Small error

Second Benckmark in scenario 1

Benchmark: Reads all the data

Scenario: 2 PC with each slave

is running as a process

Good values
Small error

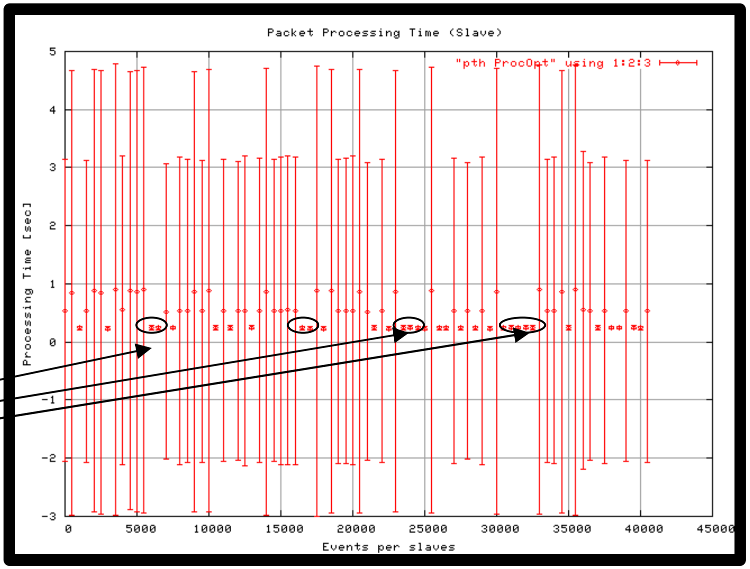

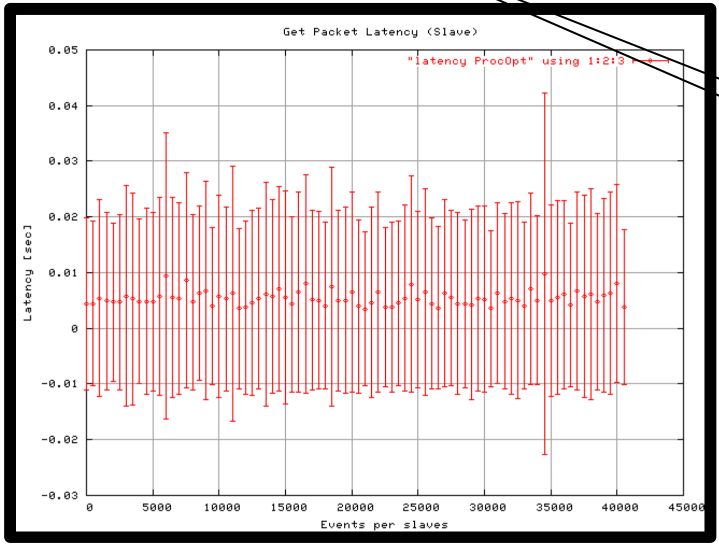

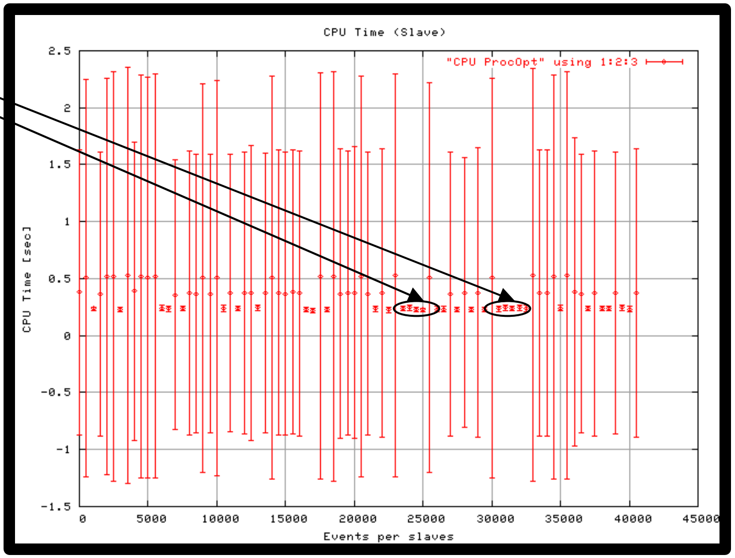

Why those huge errors?

How things **should** be

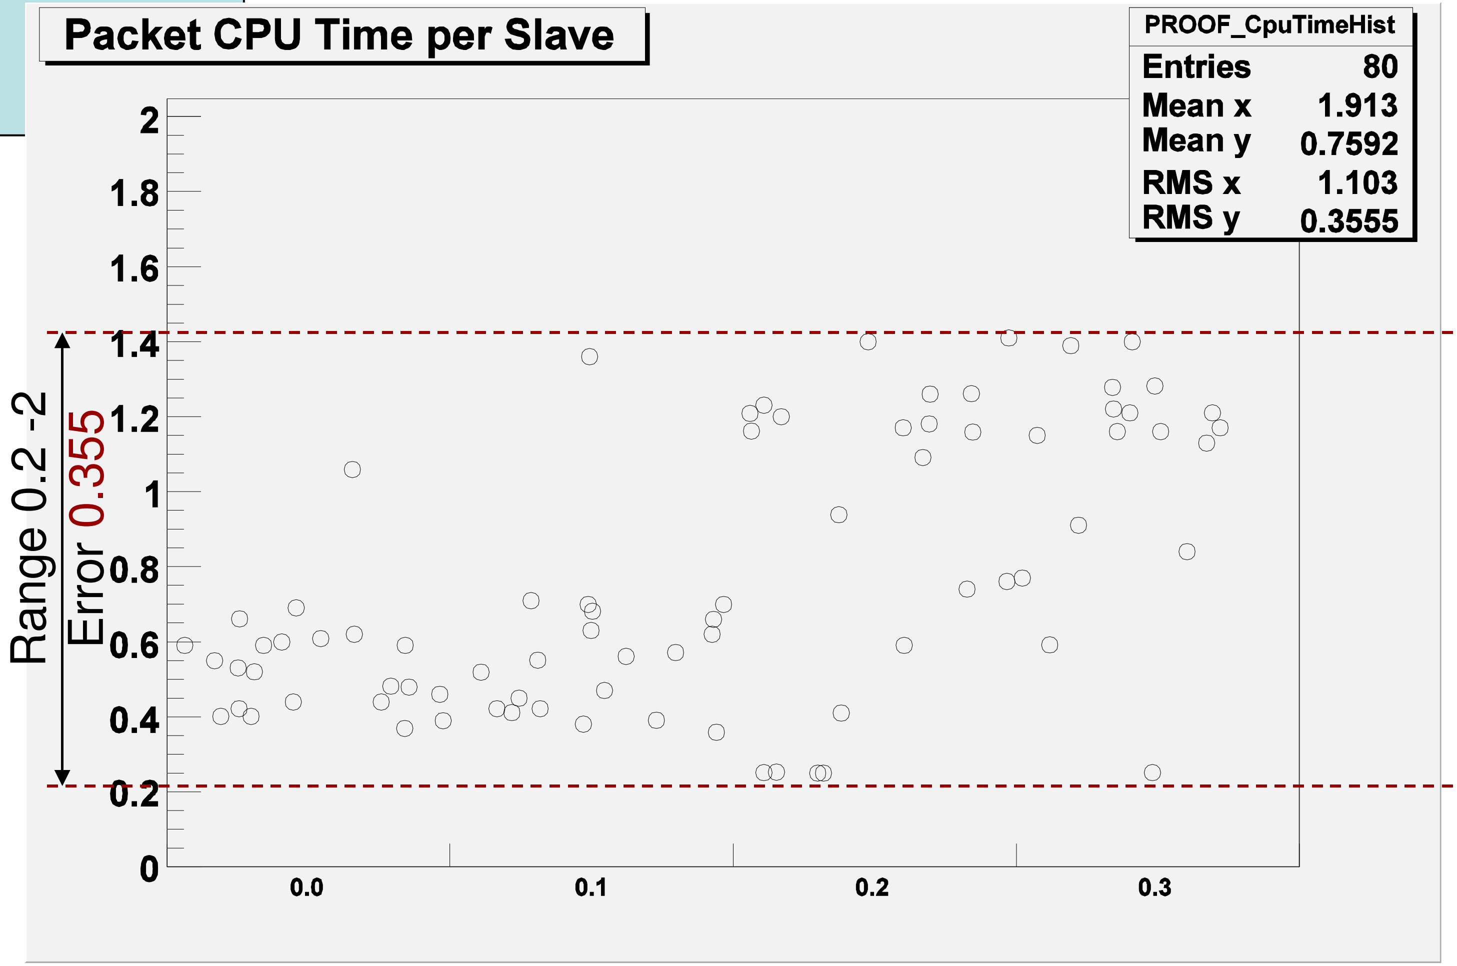

How things **are**

Why those huge errors?

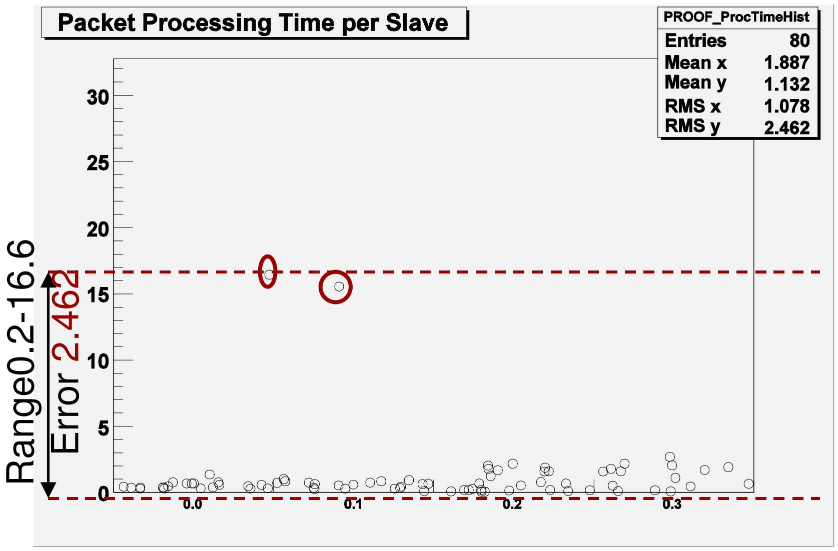

Range0.2-16.6
Error 2.462

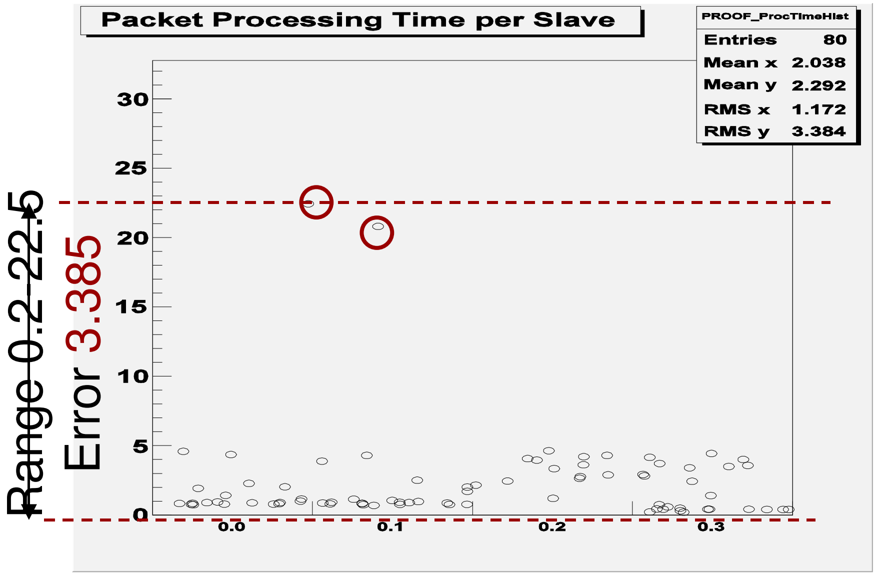

Range 0.2-22.5
Error 3.385

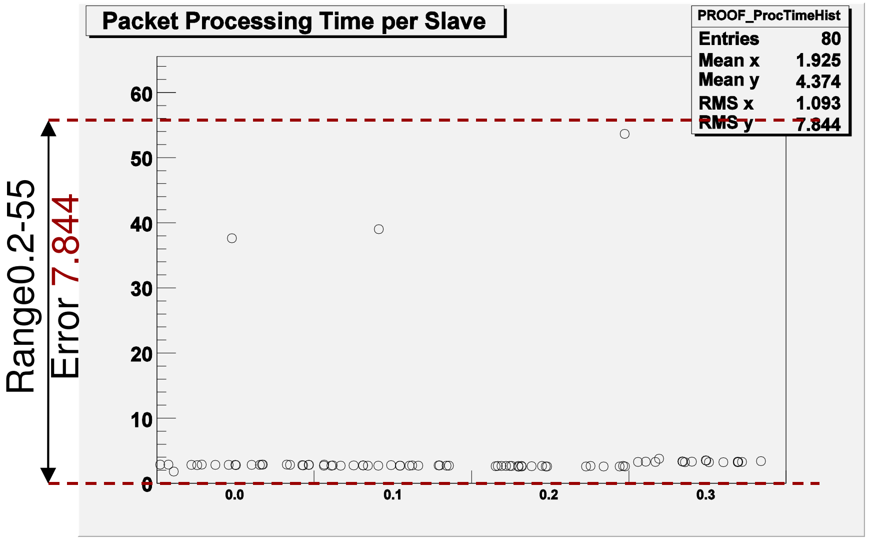

Range0.2-55
Error 7.844

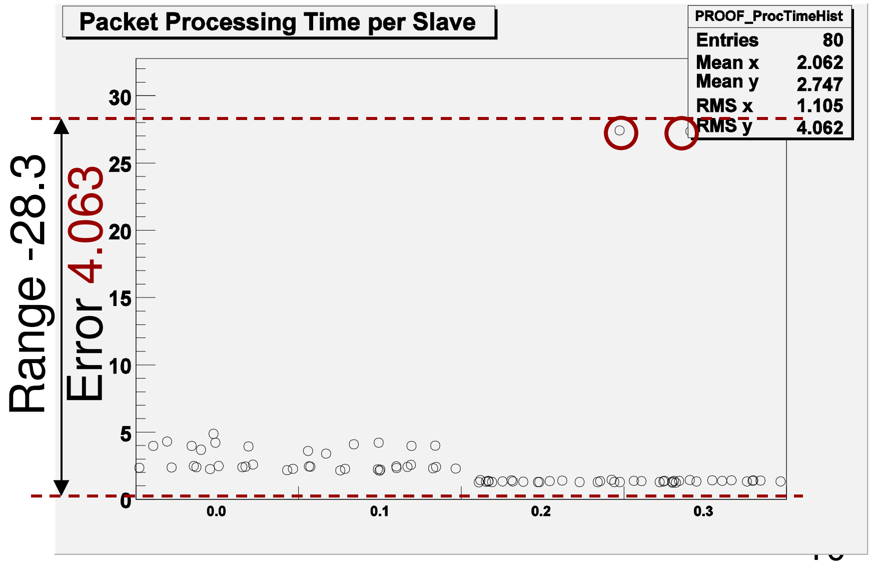

Range -28.3
Error 4.063

First benchmark in scenario 2(Xen)

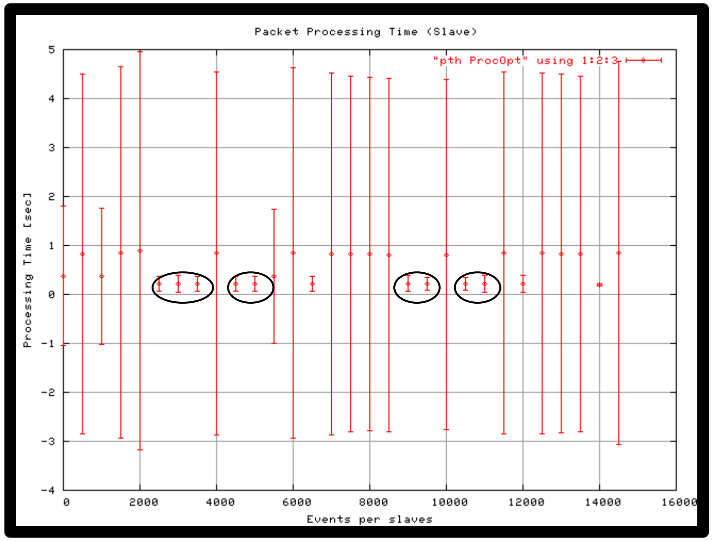

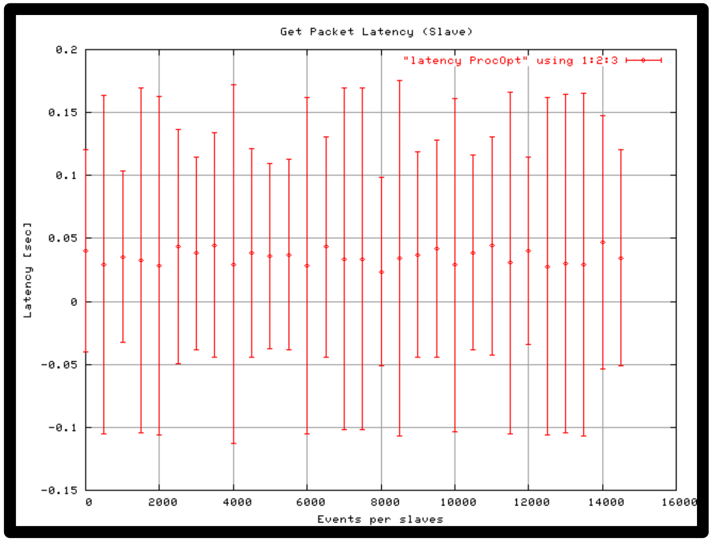

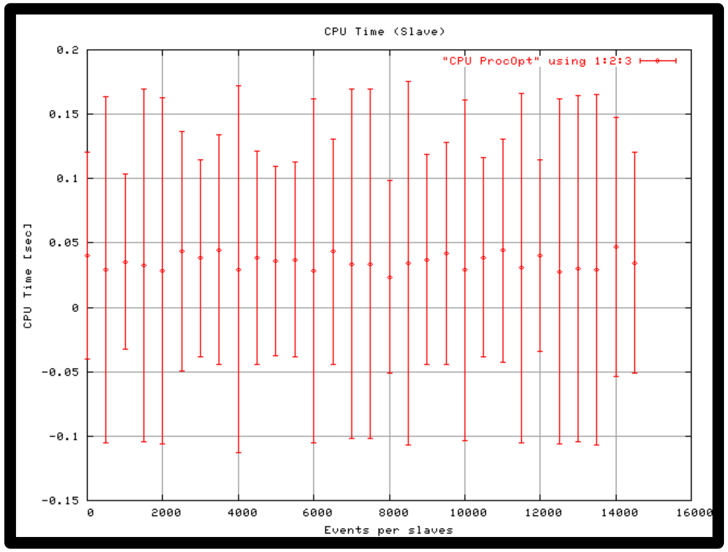

Benchmark: Reads only 25% of the data

Scenario: 2 PC with each slave runs on a guest domain

Second Benchmark in scenario 2

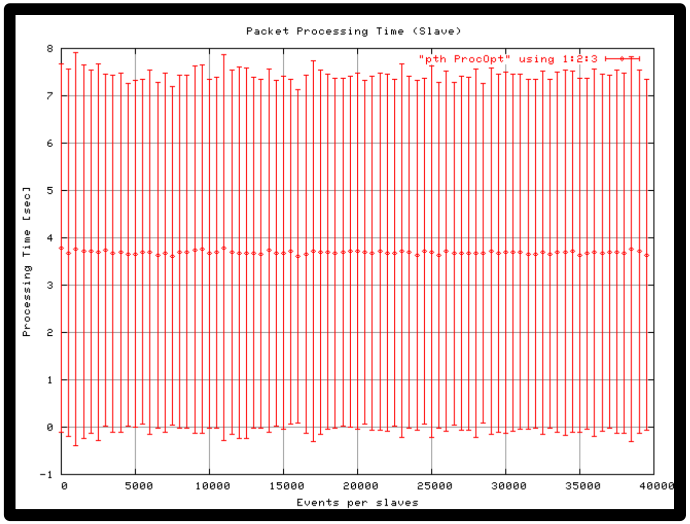

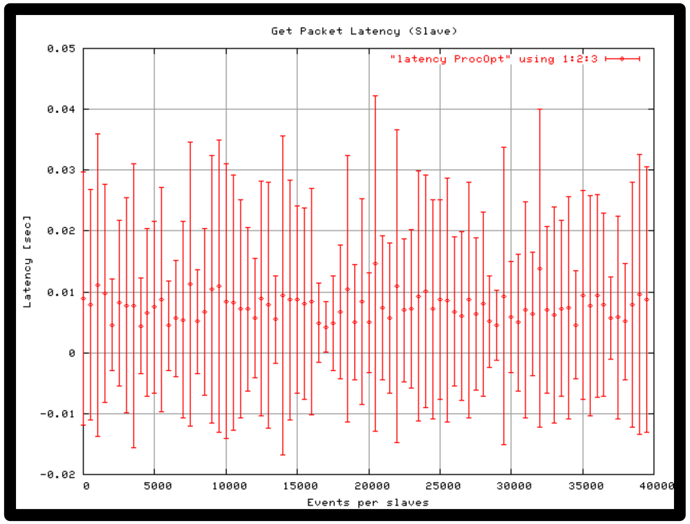

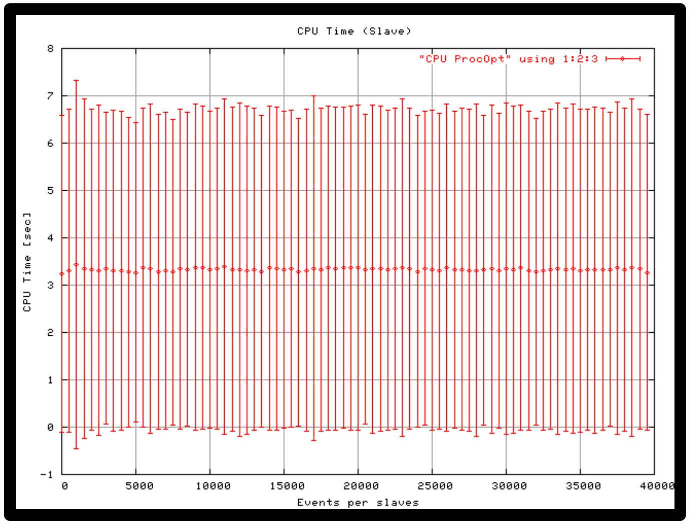

Graphical comparison Xen-PCs

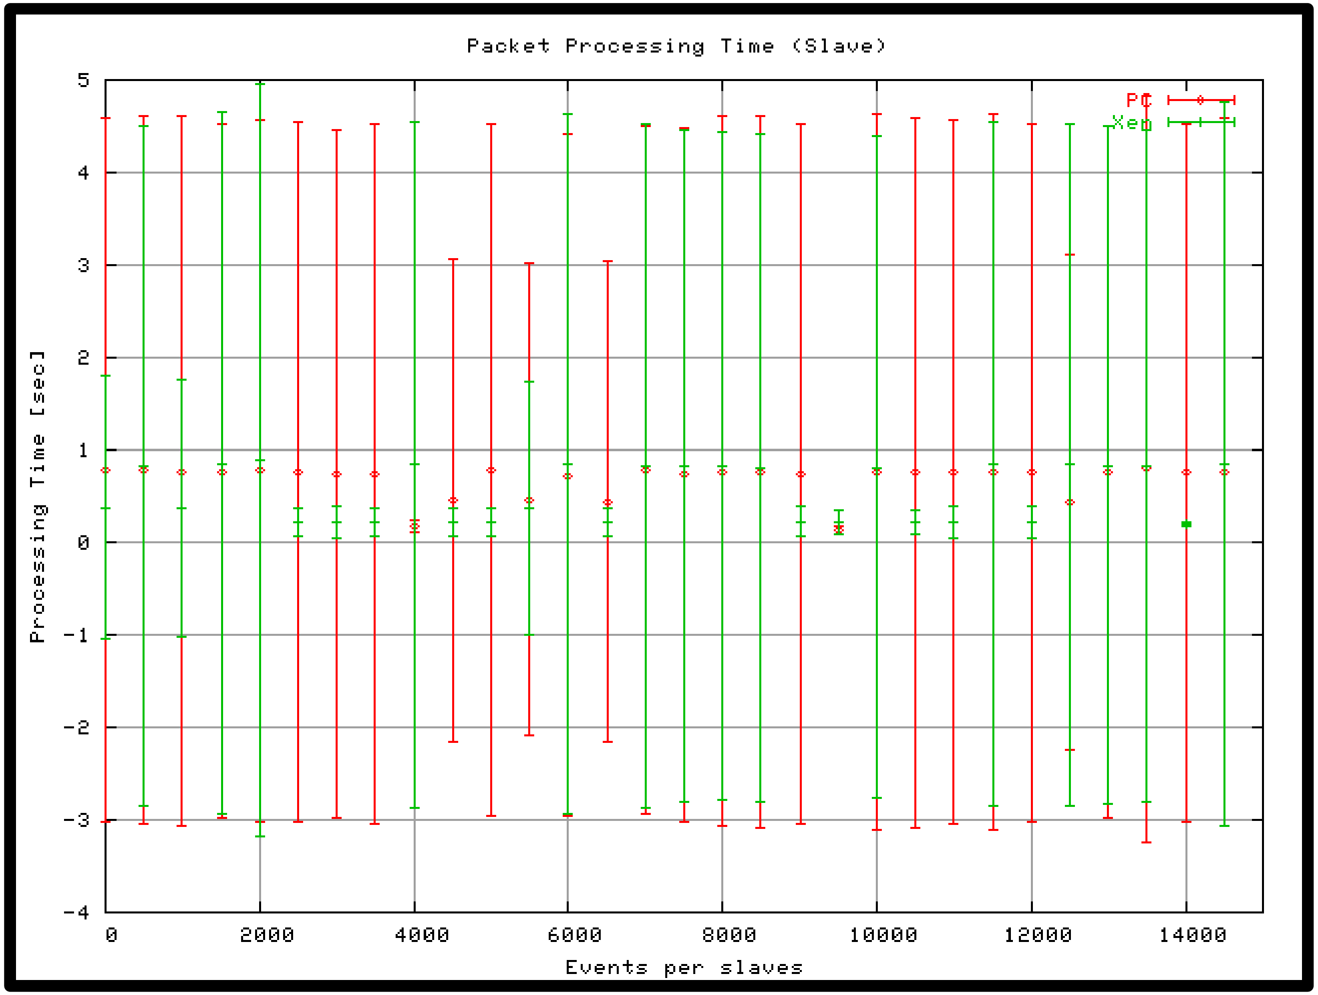

Graphical comparison Xen-PCs

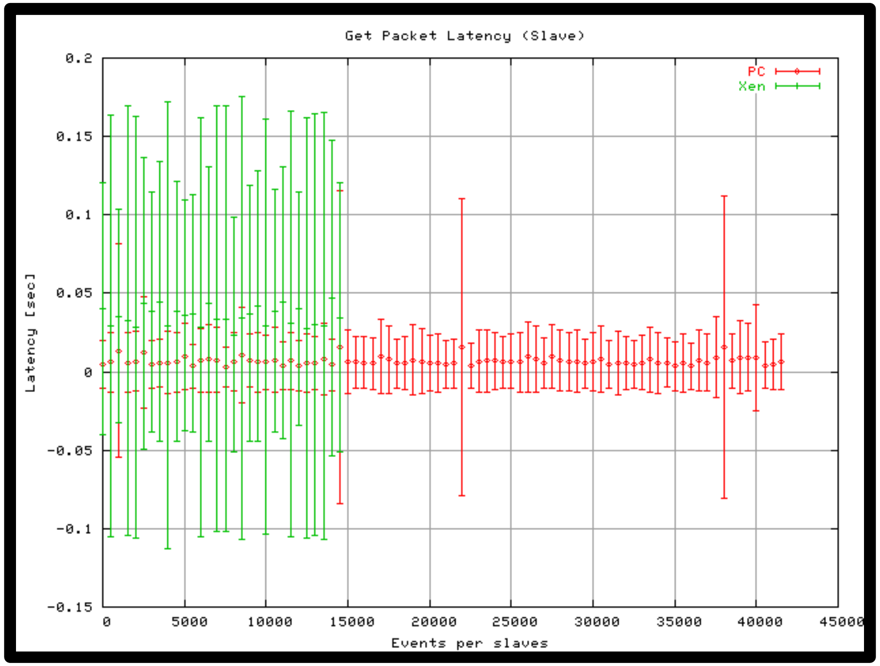

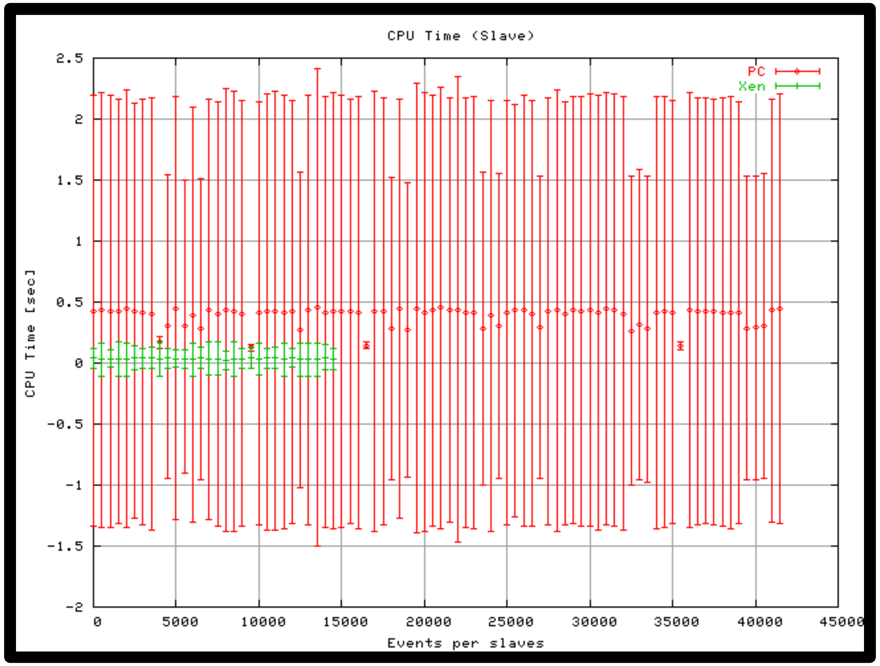

Comparison

Unreliable comparison because of the errors

		Latency	Proc	CPU
PROCESS	Opt	0.0060823	0.767716	0.418625
	Proc	0.00664648	0.771708	0.771708
XEN	Opt	0.0341873	0.810273	0.386286
	Proc	0.0444994	0.69751	0.415625

Conclusions

- As the error in the data is significant the result (CPU time) is not relevant – no realistic information.
- In the communication , last transmitions are enormous, and introduce a lot of noice
  - Bottleneck in the master
  - Possible solution   move the master to another node

Questions, comments …