




# PROOF on the Grid



Marco MEONI

CERN - Offline Week - 18.03.2009



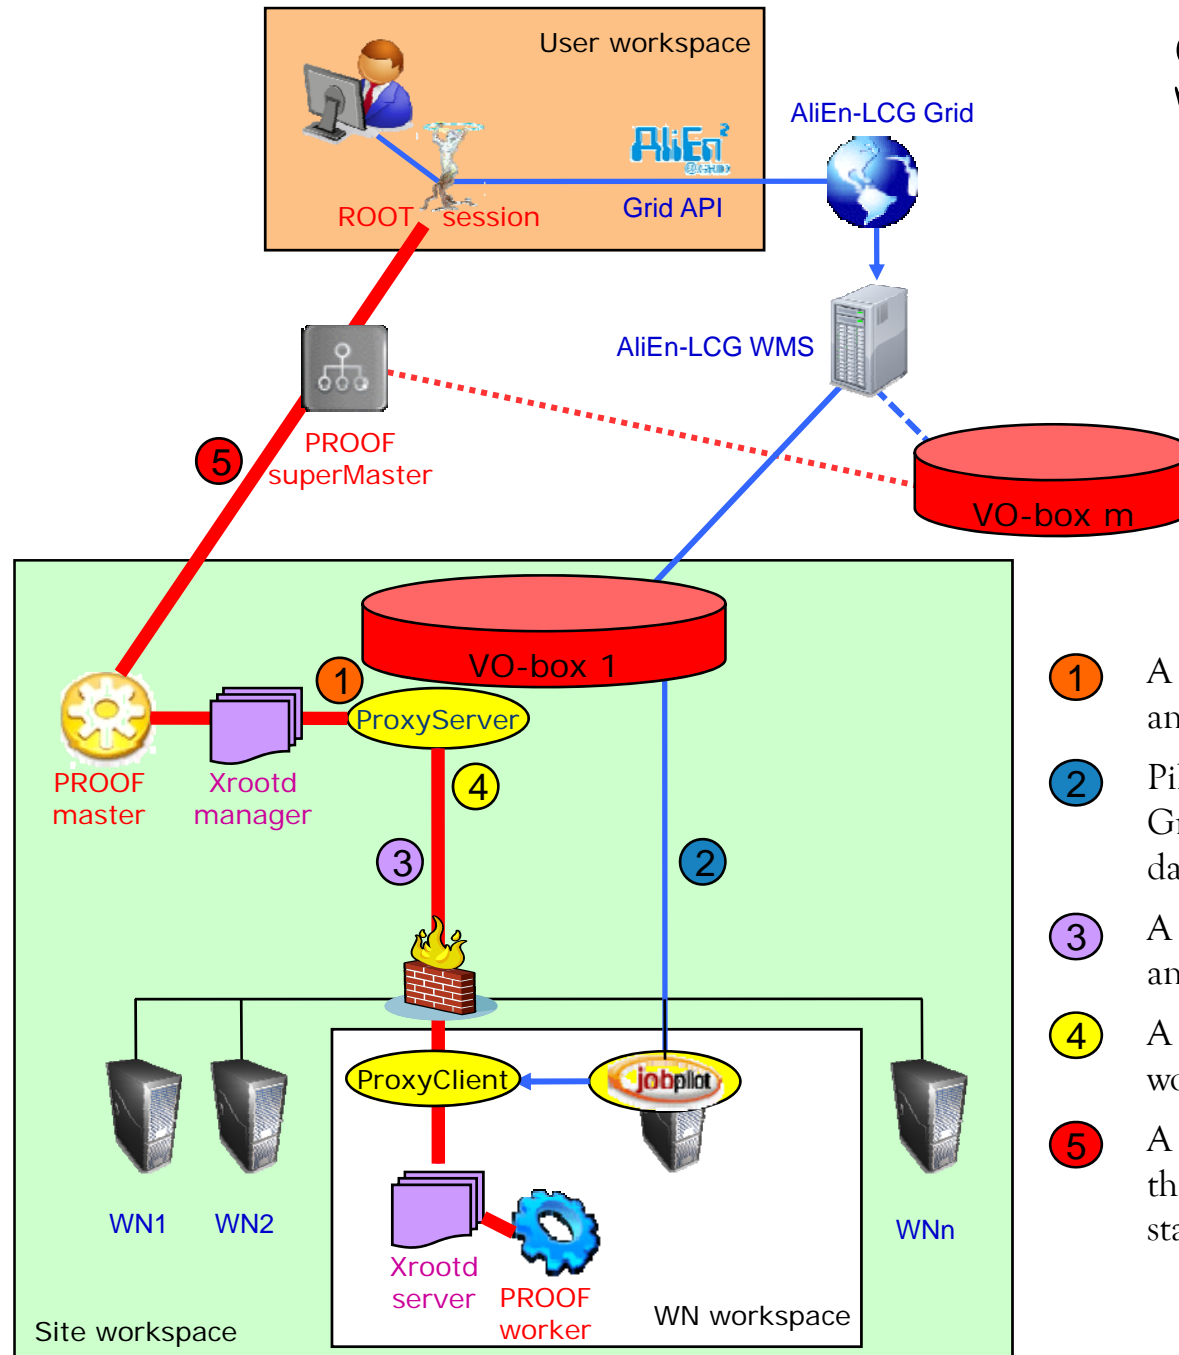
# Why?

- A cluster of the size of CAF can only hold a fraction of the yearly reconstructed and simulated data accumulated by ALICE
- It is not feasible to handle such data volume in one single computing centre
- Usage of Grid resources instead of separate hardware
- The Grid middleware knows data location and provides the classes to run Grid PROOF workers next to the data

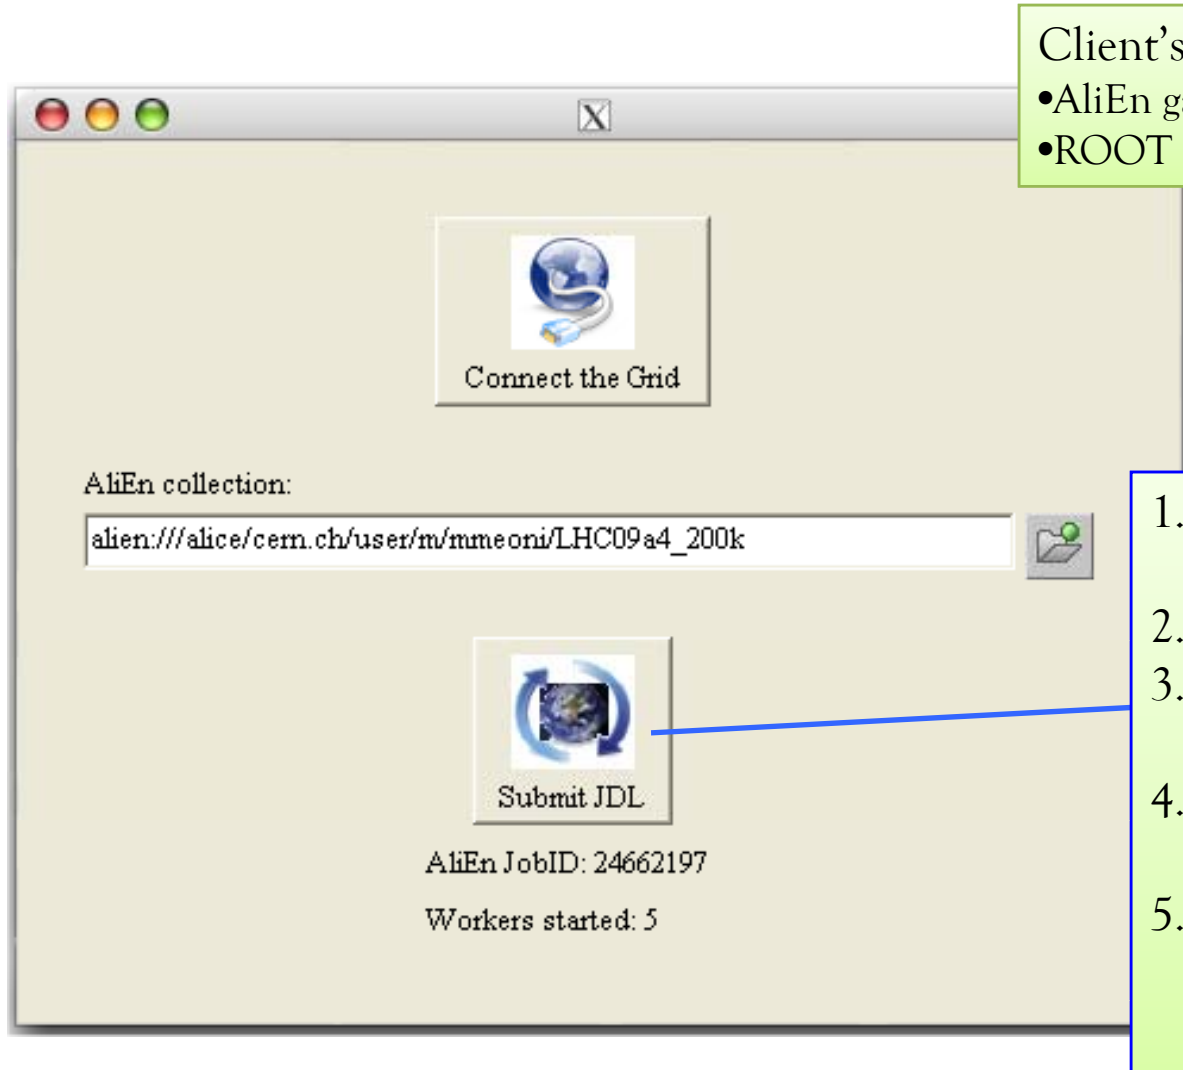
# Goals

- Cluster connectivity
  - interconnection of Grid sites running local PROOF clusters
- Tasks and Data co-location
  - the user tasks must be executed in the computing centre hosting the data
- Interactivity
  - the Grid latency must be hidden
- Protected access
  - client machines must be able to connect remote workers firewall-protected
- Workers Dynamic Allocation
  - on a Grid environment data servers may become available or drop out at any point in time

# Schema



# User Interface



Client's requirements

- AliEn gapi
- ROOT

Connect the Grid

AliEn collection:

alien:///alice/cern.ch/user/m/mmeoni/LHC09a4\_200k

Submit JDL

AliEn JobID: 24662197

Workers started: 5

1. Conversion of the AliEn collection into a text file list
2. JDL automatic creation
3. AliEn job submission (proxy already created)
4. TChain creation for analysis with PROOF
5. Options to specify ROOT version, output dir, SE requirements

# Current Prototype @ CERN

VoBox

AliEn WMS

User box

```
[voalice06] ~> netstat -vatn | grep LIST
tcp    0    0  0.0.0.0:20000  0.0.0.0:*    LISTEN
tcp    0    0  0.0.0.0:1093   0.0.0.0:*    LISTEN
tcp    0    0  0.0.0.0:1094   0.0.0.0:*    LISTEN

[voalice06] ~> less proof.conf
master voalice06.cern.ch
```

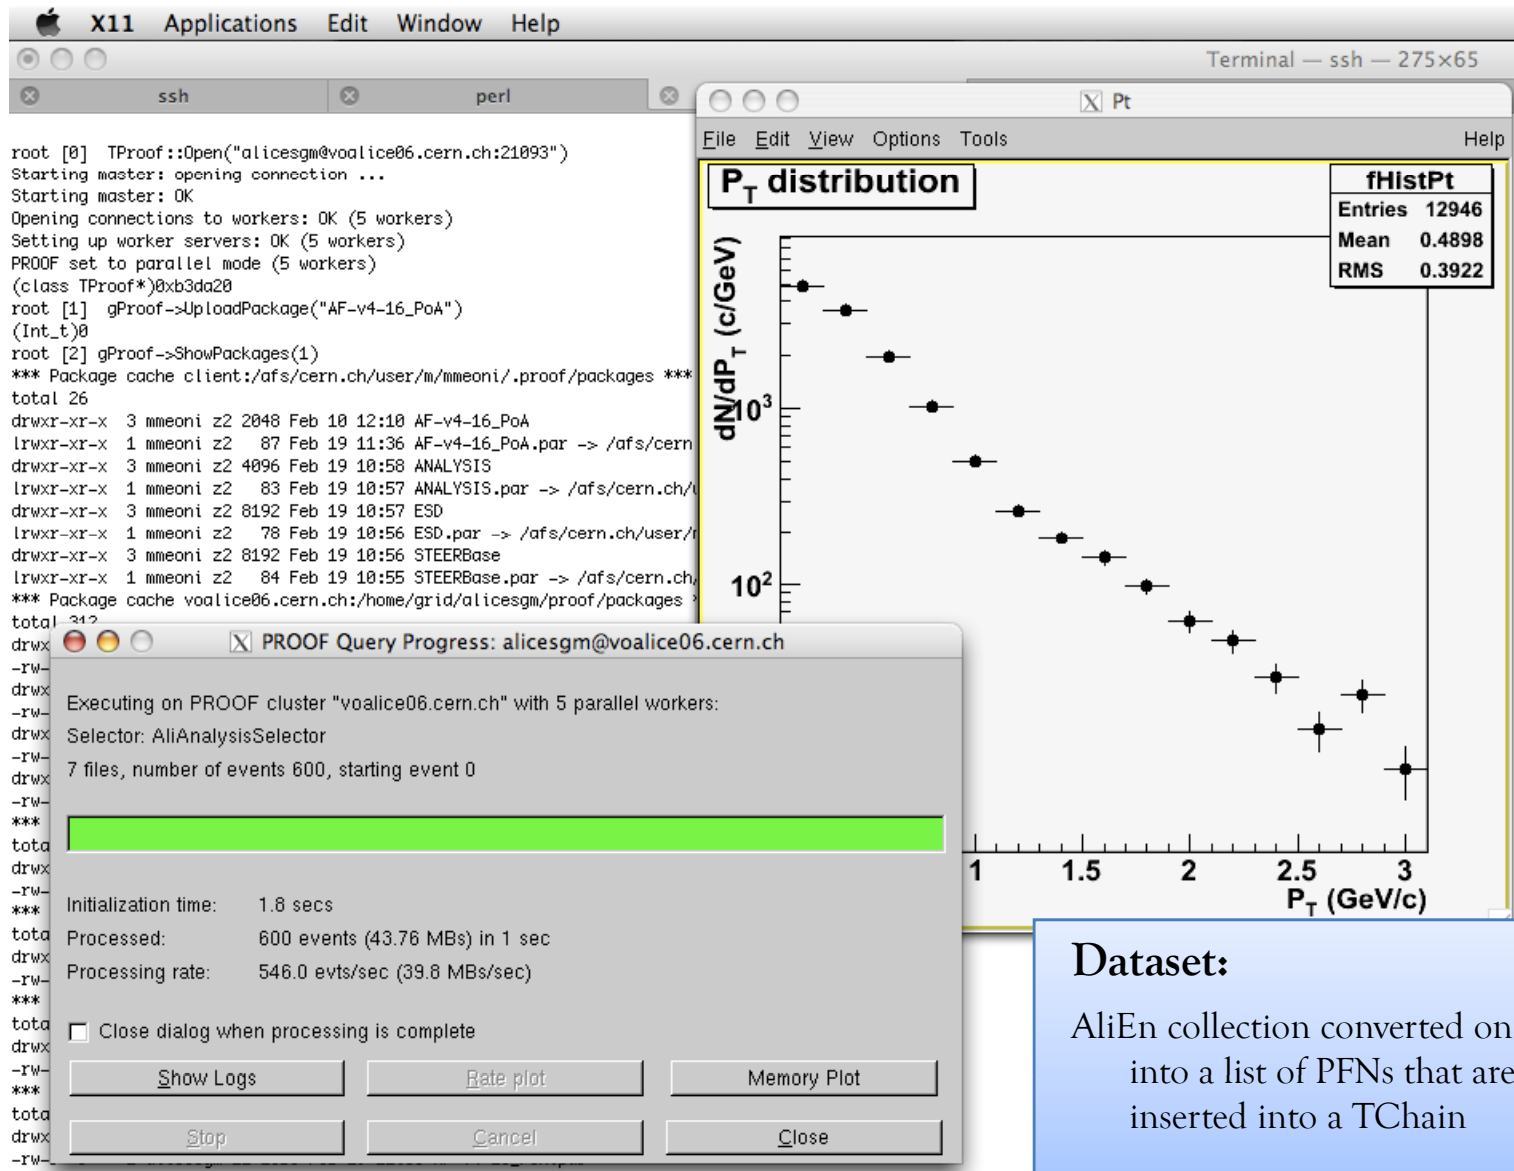
```
{ Executable="/bin/pilotJob.sh";
  Packages={"VO_ALICE@ROOT:v5-22-00a",...};
  Jobtag = {"comment: PROOF on the Grid"};
  InputFile={"LF:~/proofGrid/xpd.cf",
             "LF:~/proofGrid/proofagent.cfg.xml"};
  InputDataCollection="LF:~/LHC09a4_200k,nodownload";
  Split="file";
  OutputDir="~/temp/#alien_split%03i#";
```

```
[aliendb06c.cern.ch:3307] ~/ > ps -s -id 23180132
mmeoni    22792005  RS           ~/bin/pilotJob.sh
mmeoni    -23180133  R    00:00:00  ~/bin/pilotJob.sh
mmeoni    -23180134  R    00:00:00  ~/bin/pilotJob.sh
mmeoni    -23180135  R    00:00:01  ~/bin/pilotJob.sh
mmeoni    -23180136  R    00:00:01  ~/bin/pilotJob.sh
mmeoni    -23180136  R    00:00:01  ~/bin/pilotJob.sh
```

```
[voalice06] ~> less proof.conf
master voalice06.cern.ch
worker alicesgm@localhost port=20001 image=1
worker alicesgm@localhost port=20002 image=2
worker alicesgm@localhost port=20003 image=3
worker alicesgm@localhost port=20004 image=4
worker alicesgm@localhost port=20005 image=5
```

```
root [0] TProof::Open("alicesgm@voalice06.cern.ch")
Starting master: opening connection ...
Starting master: OK
Opening connections to workers: OK (5 workers)
Setting up worker servers: OK (5 workers)
PROOF set to parallel mode (5 workers)
(class TProof*)0xafa6e0
root [1]
```

# PROOF on the Grid in action



## Dataset:

AliEn collection converted on the fly into a list of PFNs that are inserted into a TChain

# ToDo

- Scratch space for file merging
- Faster file lookup and validation
- Hierarchy of PROOF masters (public superPROOF)
- Monitoring/control tools
- Scalability/Performance test of PROOF masters
- Integration of the proxies in ML
- Package binaries distribution?
- Preemptive submission of pilot jobs to pre-allocate workers



# Credits

- PROOF team and Jan for the usual good teamwork
- Patricia, Latchezar to have provided access to the CERN VoBox
- Claudiu for the availability to use NIHAM as second test site
- Anar (GSI) for the implementation of the proxies  
(<https://subversion.gsi.de/trac/dgrid>)
- Kilian, Victor (GSI) for the profitable discussion on common problems