



# Interactive Data Analysis with PROOF

## Experience at GSI

Anna Kreshuk, A.Manafov, P.Malzacher, V.Penso, K.Schwarz

# Outline

- Introduction – about GSI
- PROOF at GSI
  - Cluster organization and setup
  - User experience and issues
- PROOF on Demand
  - gLitePROOF project
  - Extensions to systems other than gLite
- Future plans

# GSI - Gesellschaft für Schwerionenforschung German National Centre for Heavy Ion Research



~ 1000 employees  
~ 1000 guest scientists

Budget: ~ 95 Mio Euro

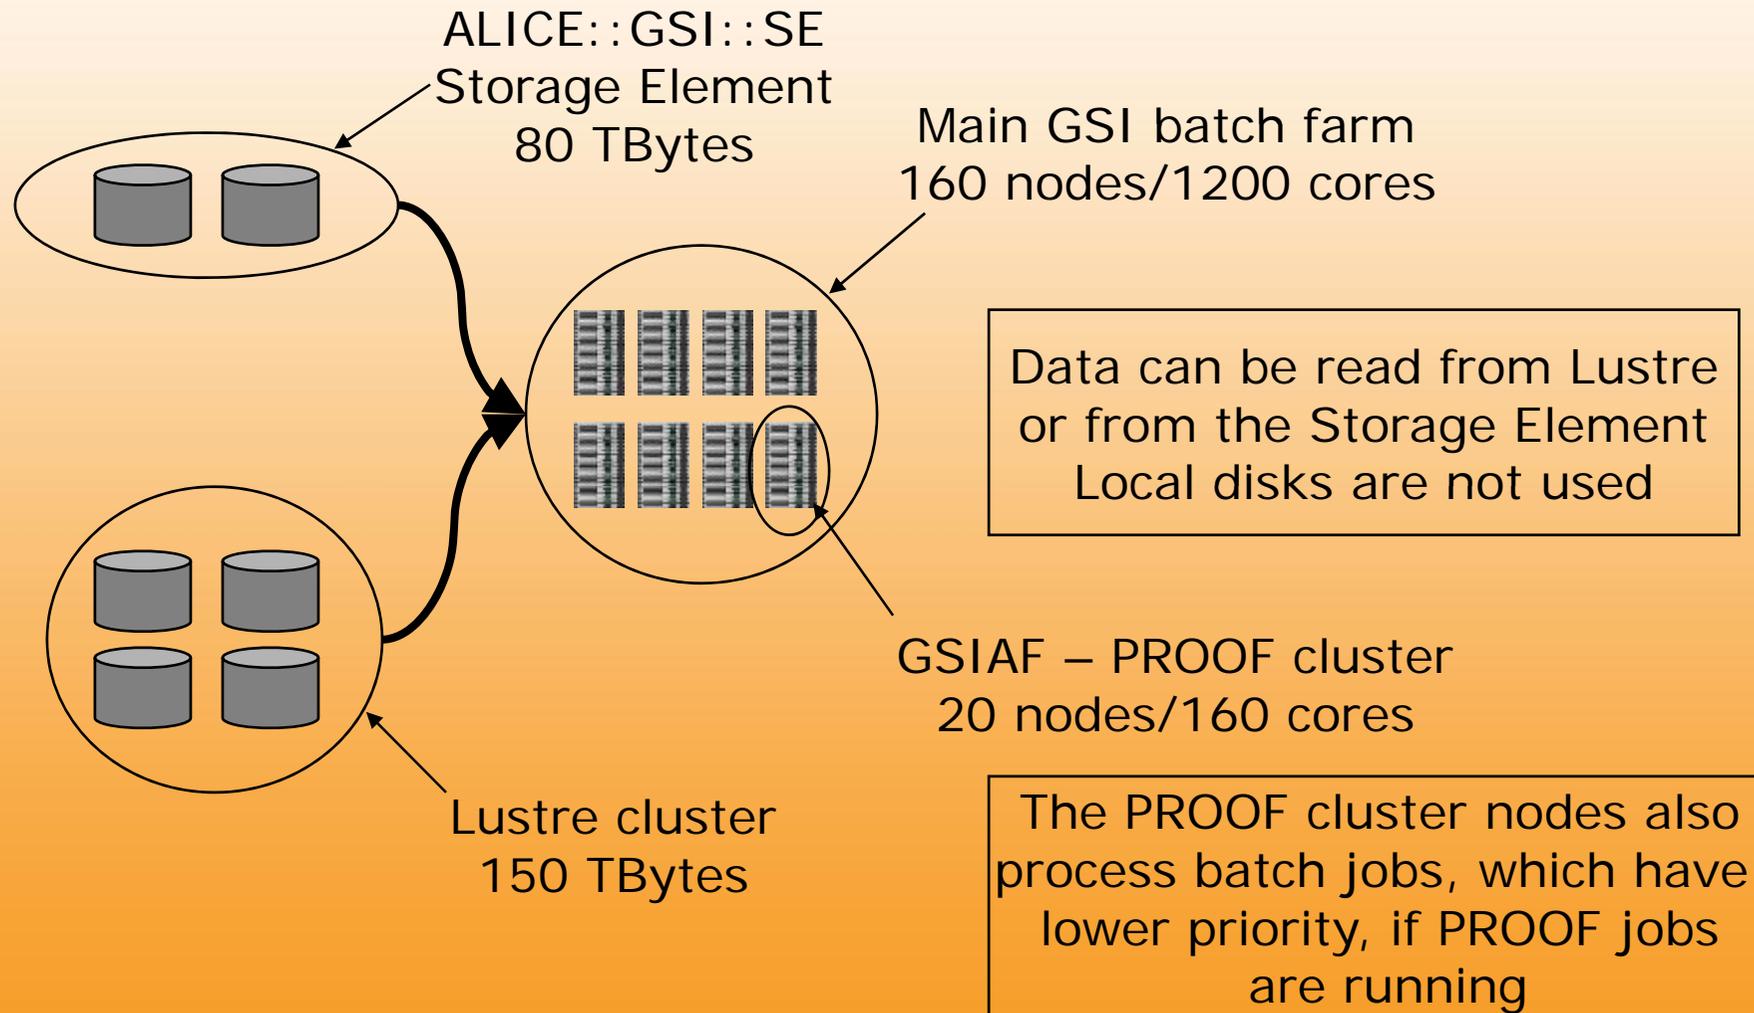


4/11/2008

ACAT 2008

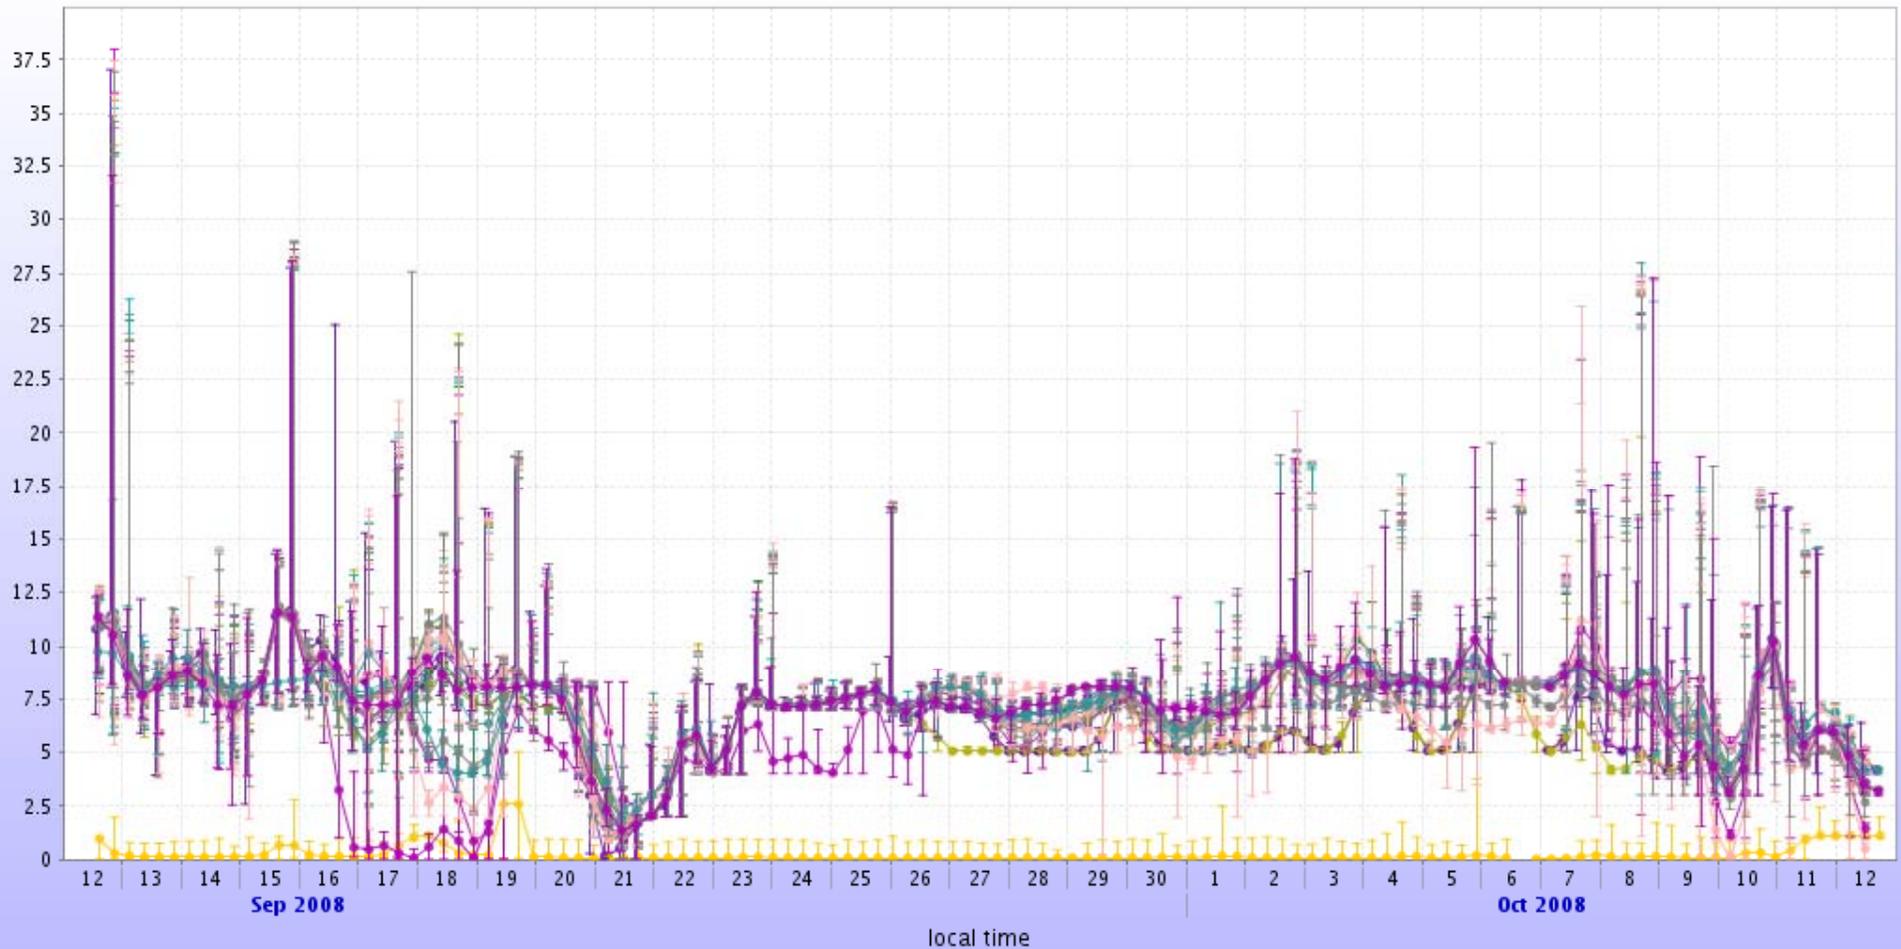
3

# GSI AF – GSI Analysis Facility



# GSI AF – GSI Analysis Facility

History of load1



- lxb341.gsi.de
- lxb342.gsi.de
- lxb343.gsi.de
- lxb344.gsi.de
- lxb346.gsi.de
- lxb347.gsi.de
- lxb348.gsi.de
- lxb349.gsi.de
- lxb350.gsi.de
- lxb351.gsi.de
- lxb352.gsi.de
- lxb353.gsi.de
- lxb354.gsi.de
- lxb355.gsi.de
- lxb356.gsi.de
- lxb357.gsi.de
- lxb358.gsi.de
- lxb359.gsi.de
- lxb360.gsi.de
- lxb361.gsi.de
- lxgrid5.gsi.de

# PROOF – user experience

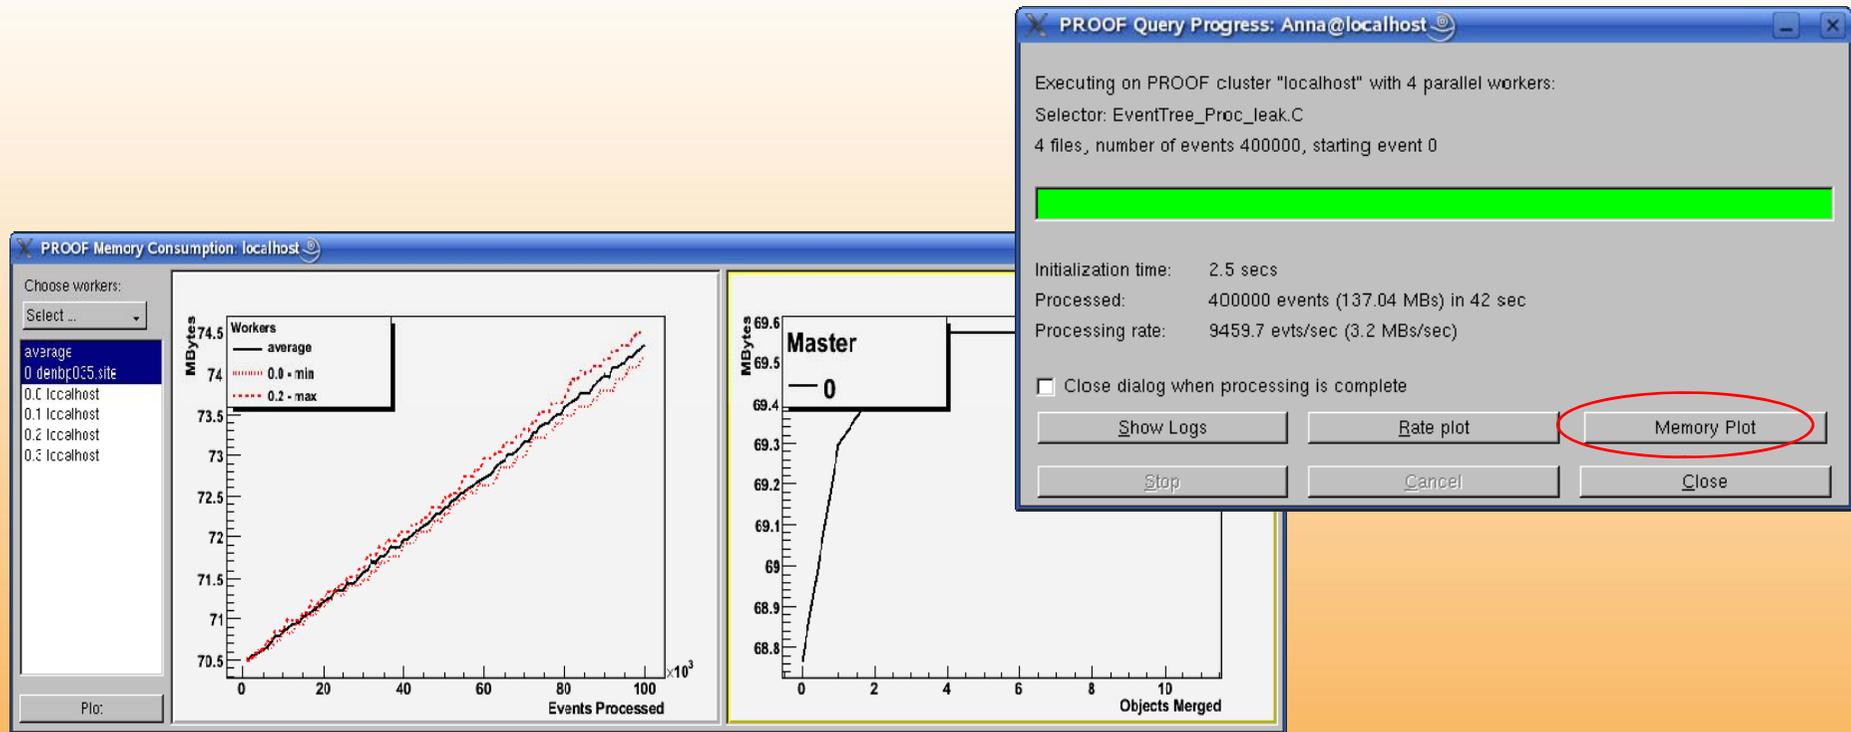
- Used heavily for code development and debugging as it provides **fast** response on **large** statistics
  - For example, ~1.4 TBytes of data are processed in ~20 minutes for a very CPU-intensive analysis
- Overall, the users are **very happy** with it
  - (almost) everything is allowed – we can still handle it with 6-8 active users
  - All machines see an NFS-mounted disk
    - users can use their own libraries
  - (almost) unlimited disk space (lustre + local disks)
    - Intermediate results at many points can be studied

# PROOF cluster - issues

- But, still there are some problems:
  - Transparency for users
    - “It runs fine locally, but crashes on PROOF, how do I find where the problem is?”
  - Fault Tolerance
    - Much progress in the last year, but still our problem #1
    - The worst is that misbehavior of one user session can kill the whole cluster
      - Happens rarely, but needs manual administrator intervention



# Improving transparency - memory



- Easily spot memory leaks on the workers and on the master
- Next step – automatic “gentle” stopping of the abusing processes

# Improving fault tolerance

- Ultimate solution – give each user his own cluster, running fully **in user space**
- Possible in GSI, because
  - any batch node can serve as PROOF node
    - no data on the local disks
  - all nodes see a shared file system with config. files
- Simple and “dirty” solution by using shell scripts
  - Submit jobs to a fast batch queue or even use lsrunc
- More complex solution, with a GUI, that allows submitting PROOF jobs both on batch and on the Grid

# gLitePROOF – PROOF on the Grid

- Developed by Anar Manafov in terms of D-Grid project
- Allows to run PROOF analysis on gLite Grid
  - A PROOF cluster with a user-defined number of workers is set up “on the fly”
  - Works with mixed type of gLite worker nodes (x86\_64, i686, etc)
  - Supports reconnections
  - Provides GUI
  - Automatic cleanup and log delivery in case of failures

# GAW - RGLite - gLitePROOF

Three modules, developed by GSI (Darmstadt) in terms of D-Grid Project AP3

## gLite Middleware

### gLite API

WMPProxy

gLite LB

Globus

LFC

### glite-api-wrapper library (GAW)

a library, which wraps parts of gLite API and adds automation and helpers to simplify access to it

### RGLite plug-in

a ROOT plug-in module, which implements the ROOT Grid interface and offers to ROOT users a possibility to use gLite middleware from within ROOT

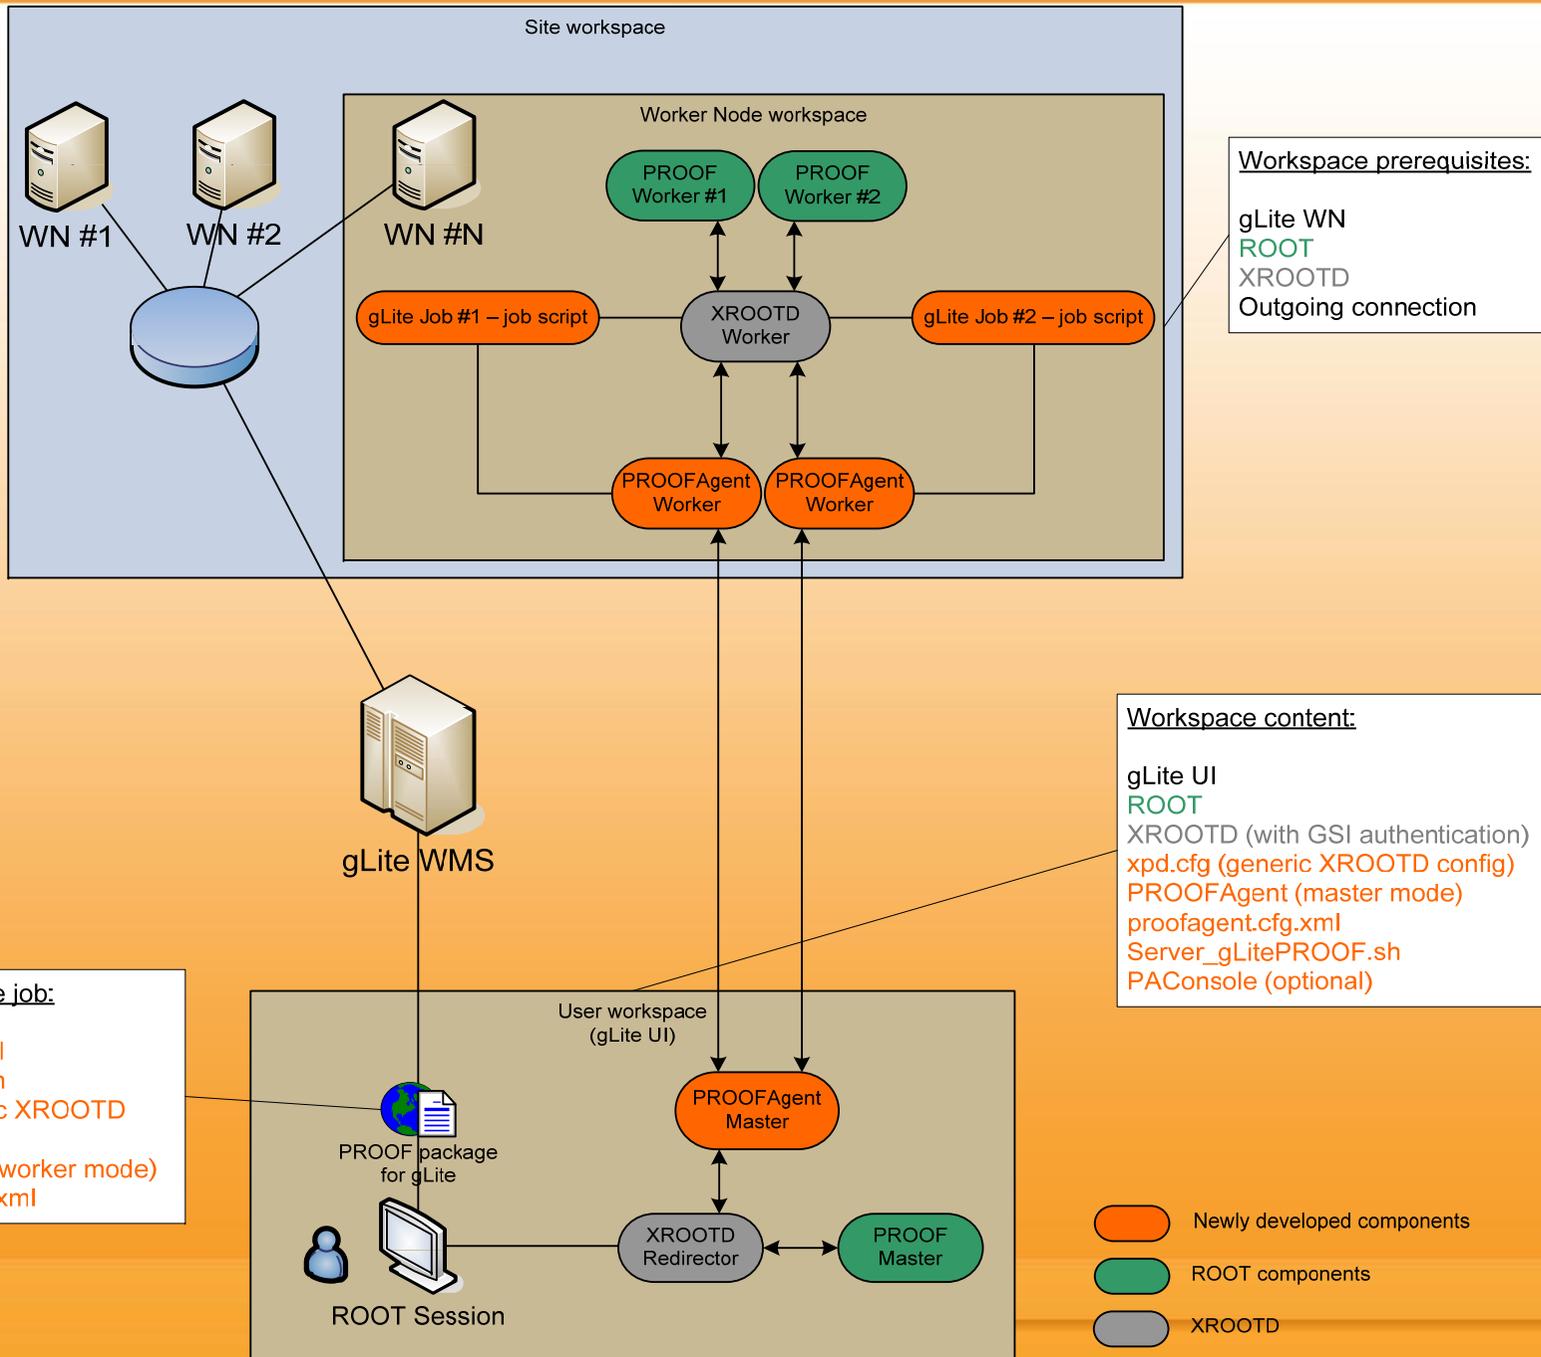
ROOT framework

### gLitePROOF

A PROOF distributed data analysis on the gLite Grid. A PROOF cluster on demand

# Components

- PROOFAgent
  - Standalone C++ application
  - Acts as a multifunctional proxy client/server
    - Provides communication layer between the PROOF master/xrootd redirector on the user machine and the PROOF workers on the Grid nodes behind a firewall
- PAConsole
  - Standalone C++ application
  - Provides a GUI to simplify user interaction with gLitePROOF
- A set of configuration files
  - Generic xrootd configuration file
  - Shell scripts for server and workers
  - A JDL file for a generic, parametric Grid job



4/11/2008

# Requirements

- User space:
  - ROOT 5.18+ (xrootd enabled)
  - BOOST 1.32
  - Qt 4.2.x (only for the GUI)
  - cmake 2.6.2 (only for the GUI)
  - gLite UI 3.1
- Worker node space:
  - Outgoing connection for GLOBUS ports (configurable)
  - gLite WMS (WMPProxy endpoint)
  - gLite WNs (at least v3.0)

# User perspective (0)

- The application is launched through a GUI PAConsole
- All configuration files can be used with default values
  - Fine tuning also possible
- As soon as a sufficient number of worker nodes have started, user can connect to PROOF from ROOT prompt as “localhost”
- gLitePROOF session is almost exactly the same as a session on a dedicated cluster
- More workers can be added on demand

# User perspective (1)

The image displays two overlapping windows of the PROOFAgent Management Console, version 1.0.5. The left window is titled "Server" and shows the status of the server. The right window is titled "Grid" and shows the configuration and status of the grid.

**Server Window:**

- Server's status:
  - XROOTD <20627>: is running
  - PROOFAgent <20629>: is running
  - PROOFAgent's version: PROOFAgent v.1.0.4
- Buttons: Start, Stop, Status
- application file name: proofagent
- protocol version: PProtocol:0.1.0
- Report bugs/comments to [A.Manafov@gsi.de](mailto:A.Manafov@gsi.de)
- PA's pid directory:
- Close button

**Grid Window:**

- JDL File:
- Endpoint:
- Submit button
- worker(s):
- 100% submitted progress bar
- Information about the last submitted job:

ID	Status
https://grid25.gsi.de:9000/1jDDPX...	Running
https://grid25.gsi.de:9000/MZpb...	Ready
https://grid25.gsi.de:9000/nIZT6...	Ready
https://grid25.gsi.de:9000/td3CA2v...	Running
https://grid25.gsi.de:9000/uECS...	Ready
https://grid25.gsi.de:9000/VAaM...	Ready
https://grid25.gsi.de:9000/KzQc...	Ready
https://grid25.gsi.de:9000/Z-923...	Ready

Close button

# Future plans

PROOF on gLite → PROOF on demand

- Make PAConsole plug-in based, with plug-ins for different job submission systems
  - gLite plug-in is ready
  - A plug-in for LSF is almost done – work in progress
- Easy setup – no centralized intervention needed
- User-defined cluster size
- Free monitoring and fair share

# Conclusion

- Our PROOF experience has been **very positive**
  - Thanks to the PROOF team for the support!
- PROOF is an excellent tool for a Tier2 center
  - Fast development on large statistics
  - No resources wasted if batch jobs run along
  - Fairly easy to setup
    - More documentation needed about medium-size simple setup
- Dynamic cluster setup in user space is the path we want to research next