

THE ALICE WORKLOAD MANAGEMENT SYSTEM: STATUS BEFORE THE REAL DATA TAKING


Stefano Bagnasco

INFN Torino

On behalf of the ALICE AliEn team

OUTLINE OF THE TALK

- Job Management model in ALICE
- Issues with gLite 3.1 WMS
- Experience with CREAM
- Conclusions



- Task Queue and Optimizers

- Central DB of jobs to be executed
- Optimizers split and arrange jobs according to input data, priority policies and/or user defined criteria

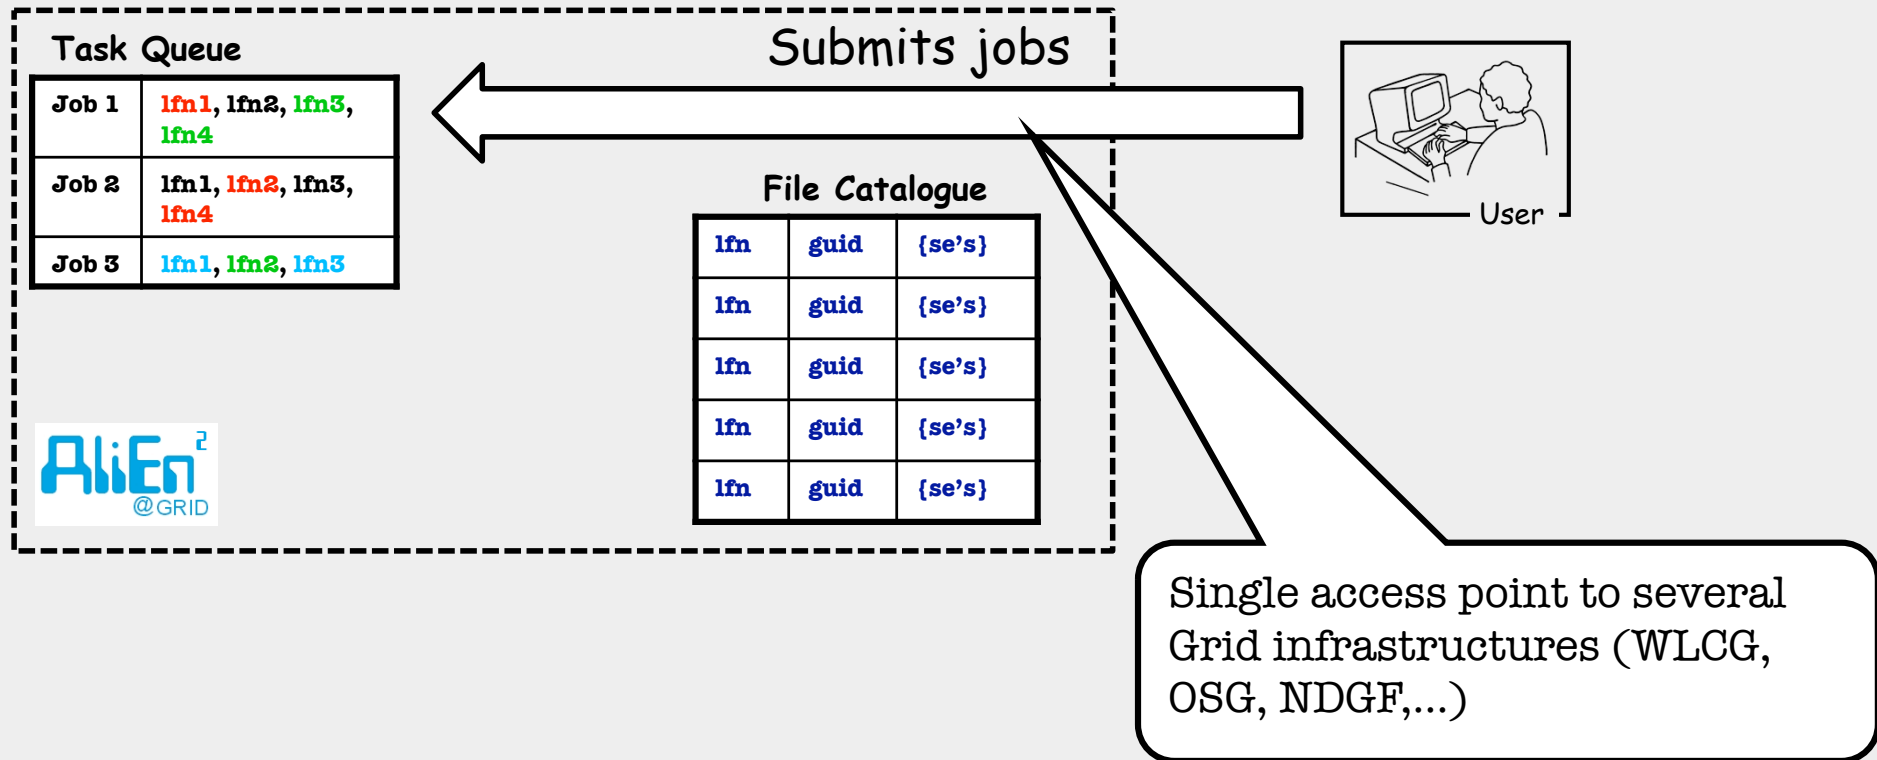
- Site VO-Box

- Thin interfaces to underlying Grid site services
- Submits JobAgents to site
- Takes care of proxy management


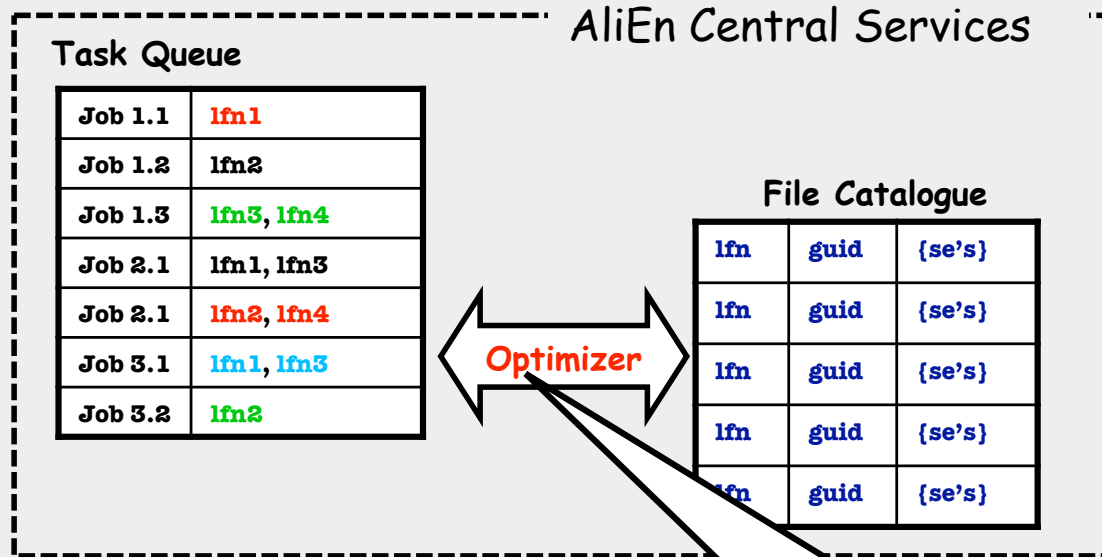
- JobAgent

- Runs on WNs, downloads payload from the TQ and executes it
- JAs create a “virtual grid” on top of existing Grid infrastructures

ALICE JOB MANAGEMENT I



ALICE JOB MANAGEMENT II



- Split master jobs into subjobs according to data locality and user-defined criteria
- Sort jobs according to priority policies
- Enforce fair share and other policies

- Task Queue and Optimizers

- Central DB of jobs to be executed
- Optimizers split and arrange jobs according to input data, priority policies and/or user defined criteria

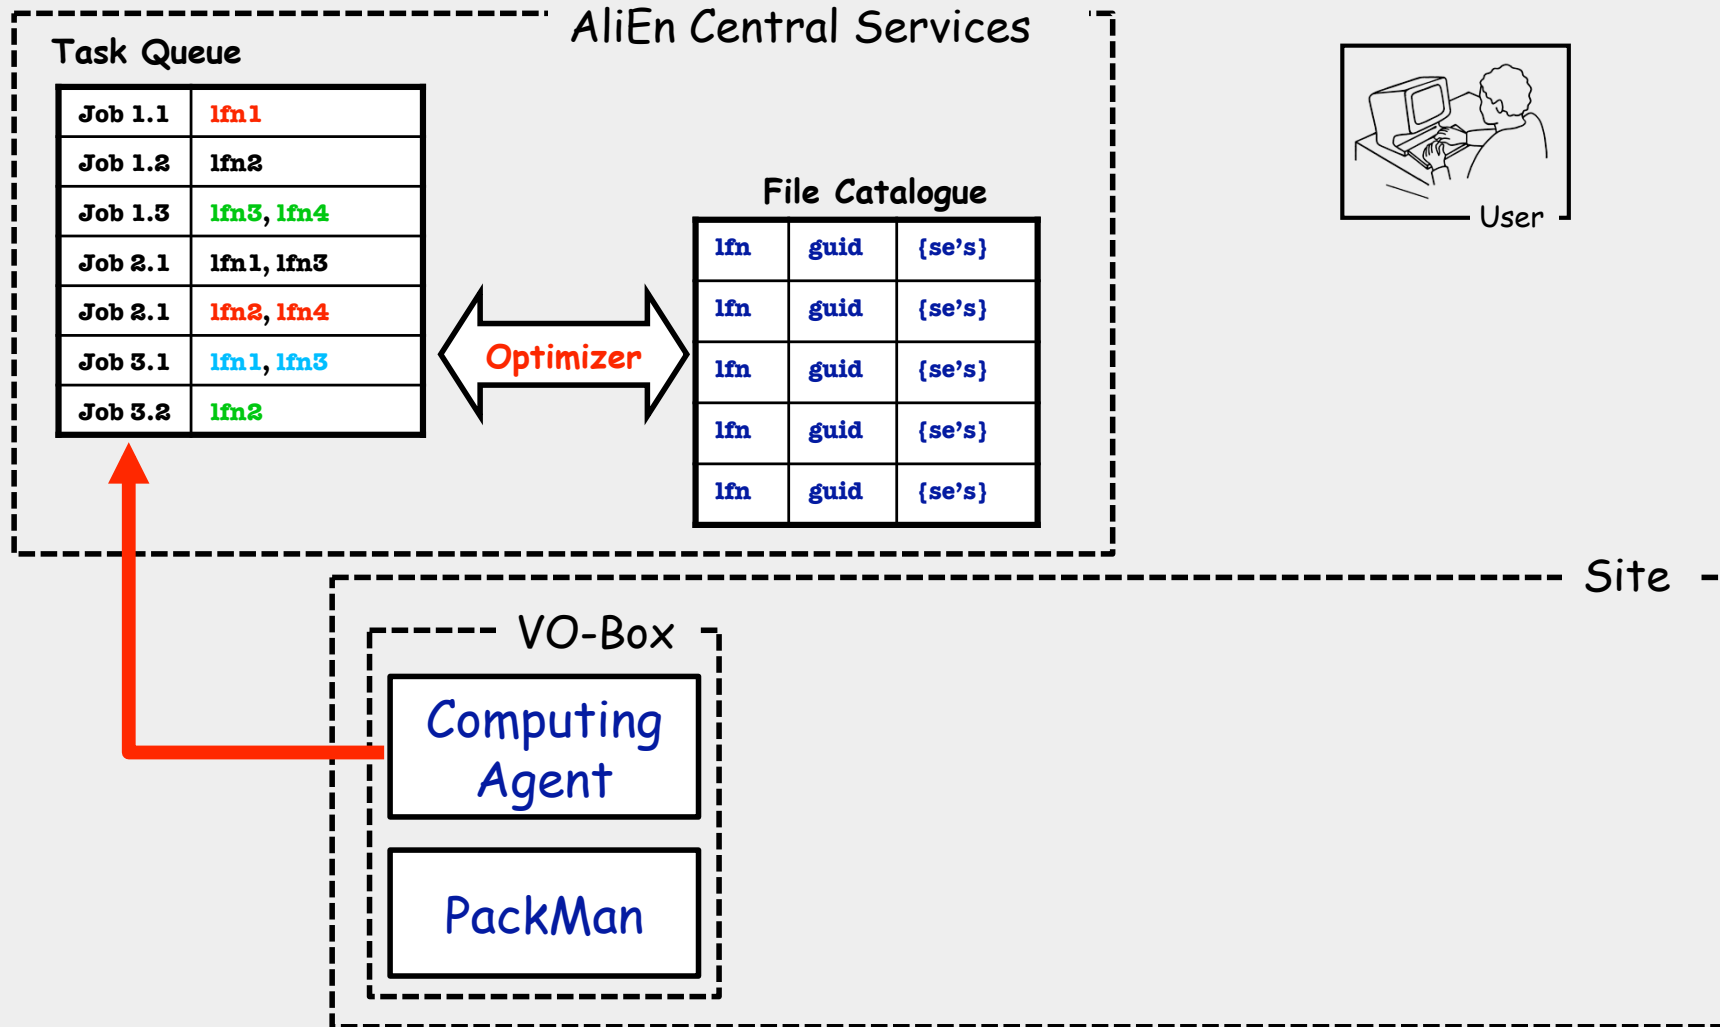
- Site VO-Box

- Thin interfaces to underlying Grid site services
- Submits JobAgents to site
- Takes care of proxy management

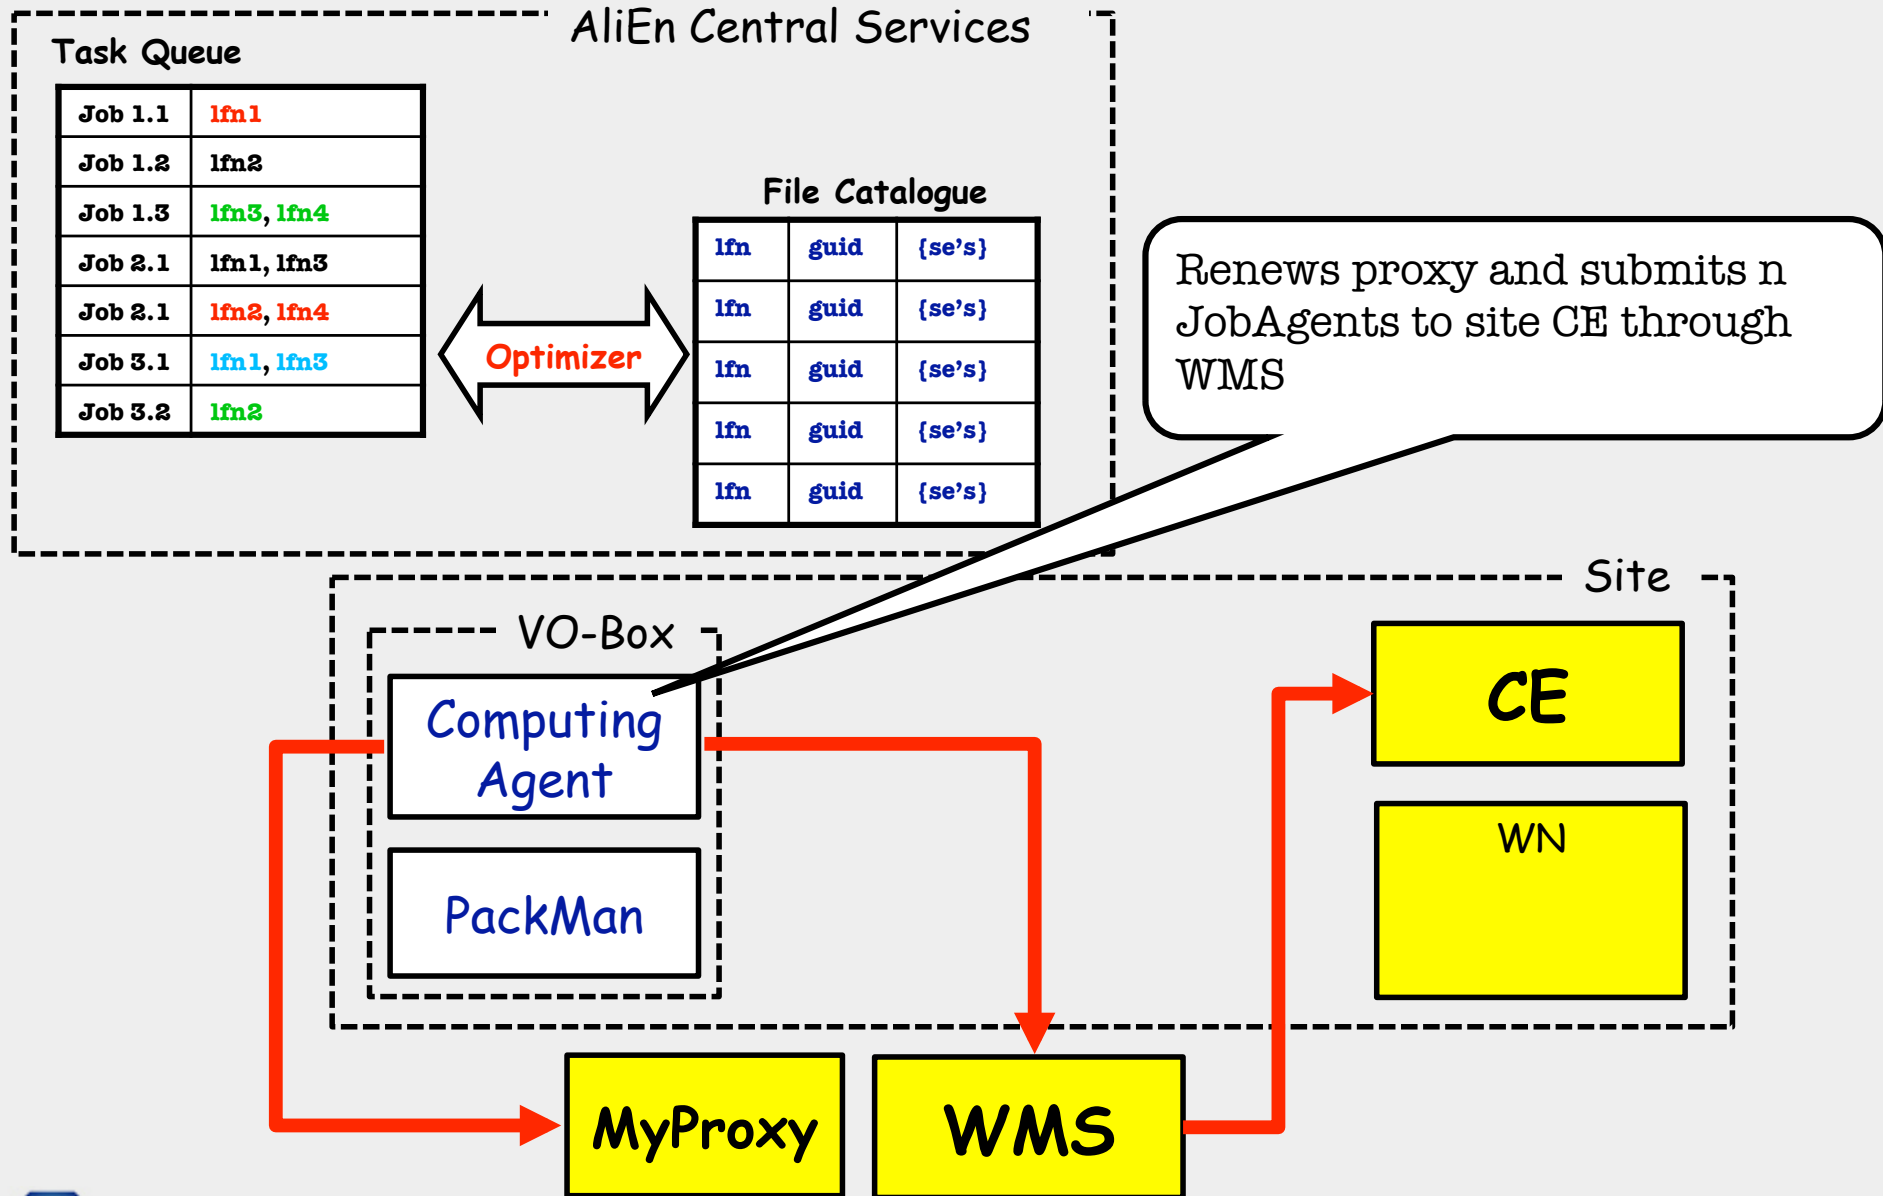
- JobAgent

- Runs on WNs, downloads payload from the TQ and executes it
- JAs create a “virtual grid” on top of existing Grid infrastructures

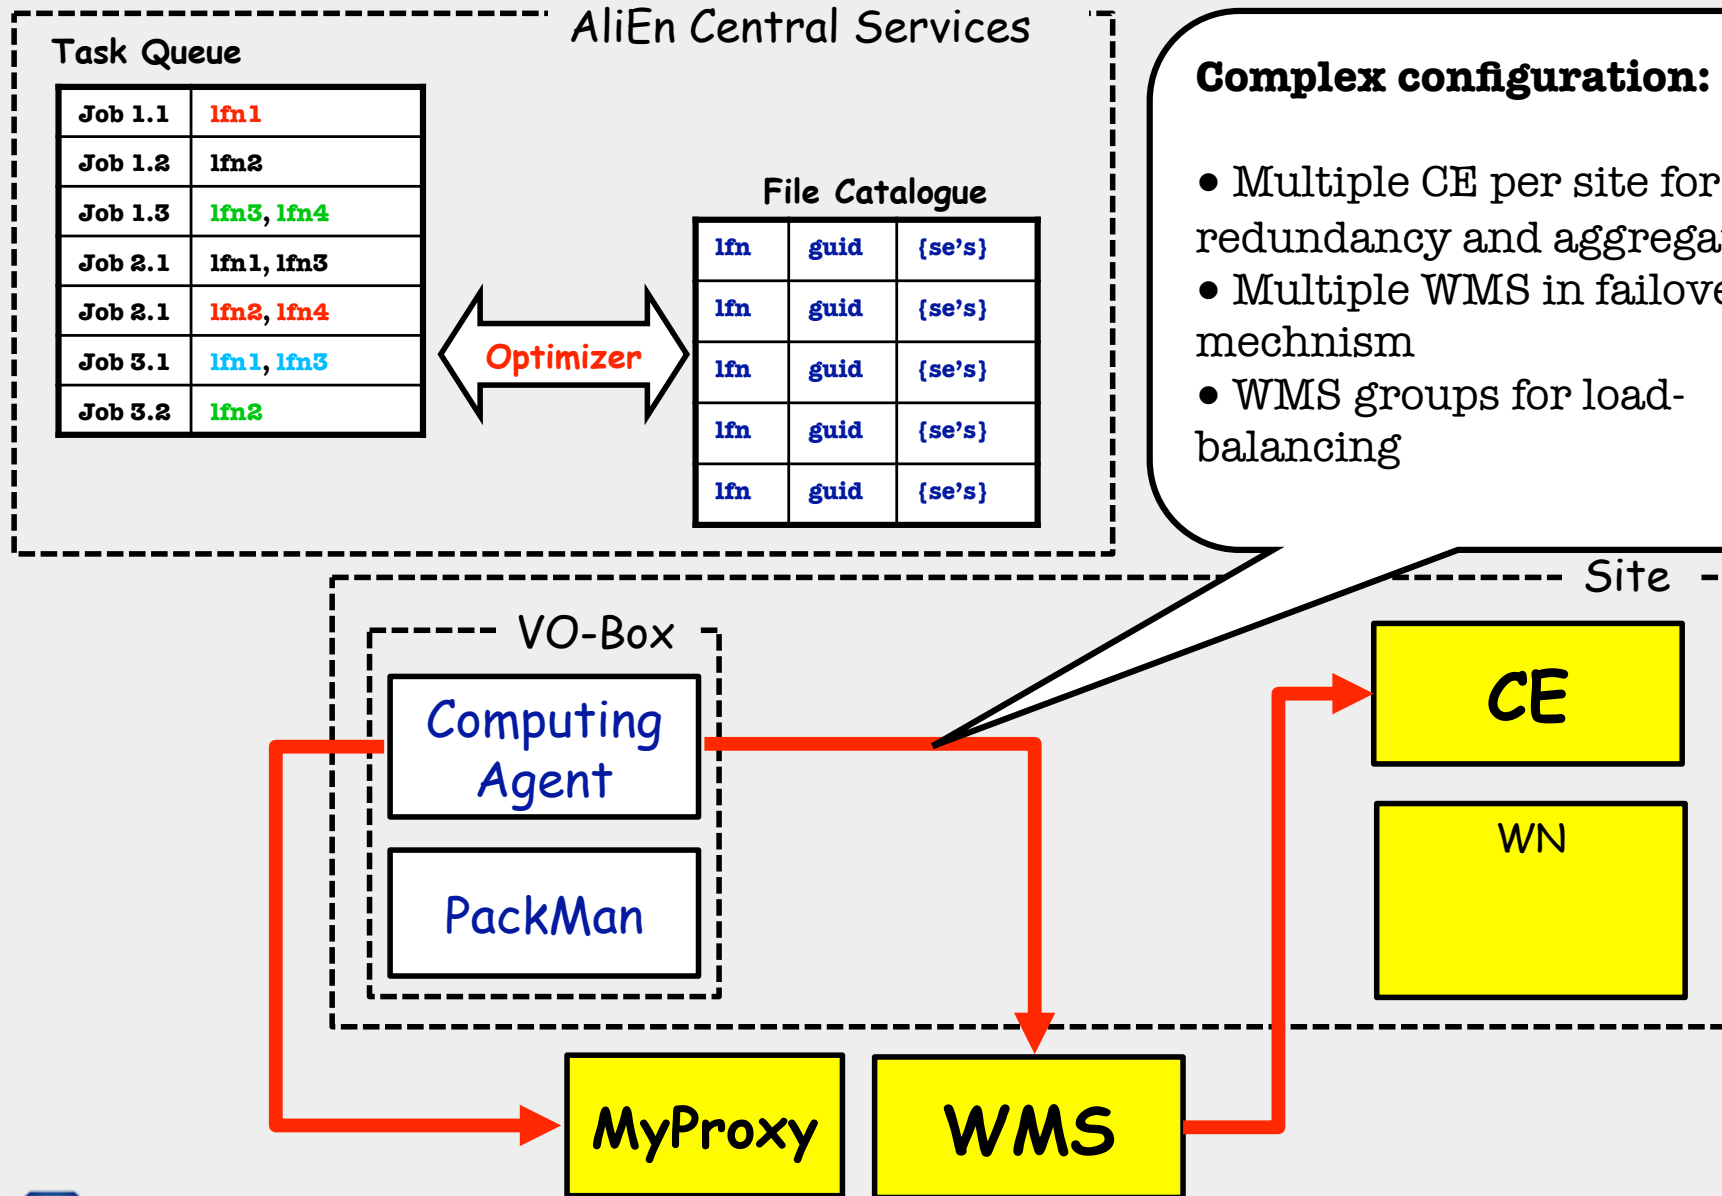
ALICE JOB MANAGEMENT III



ALICE JOB MANAGEMENT IV

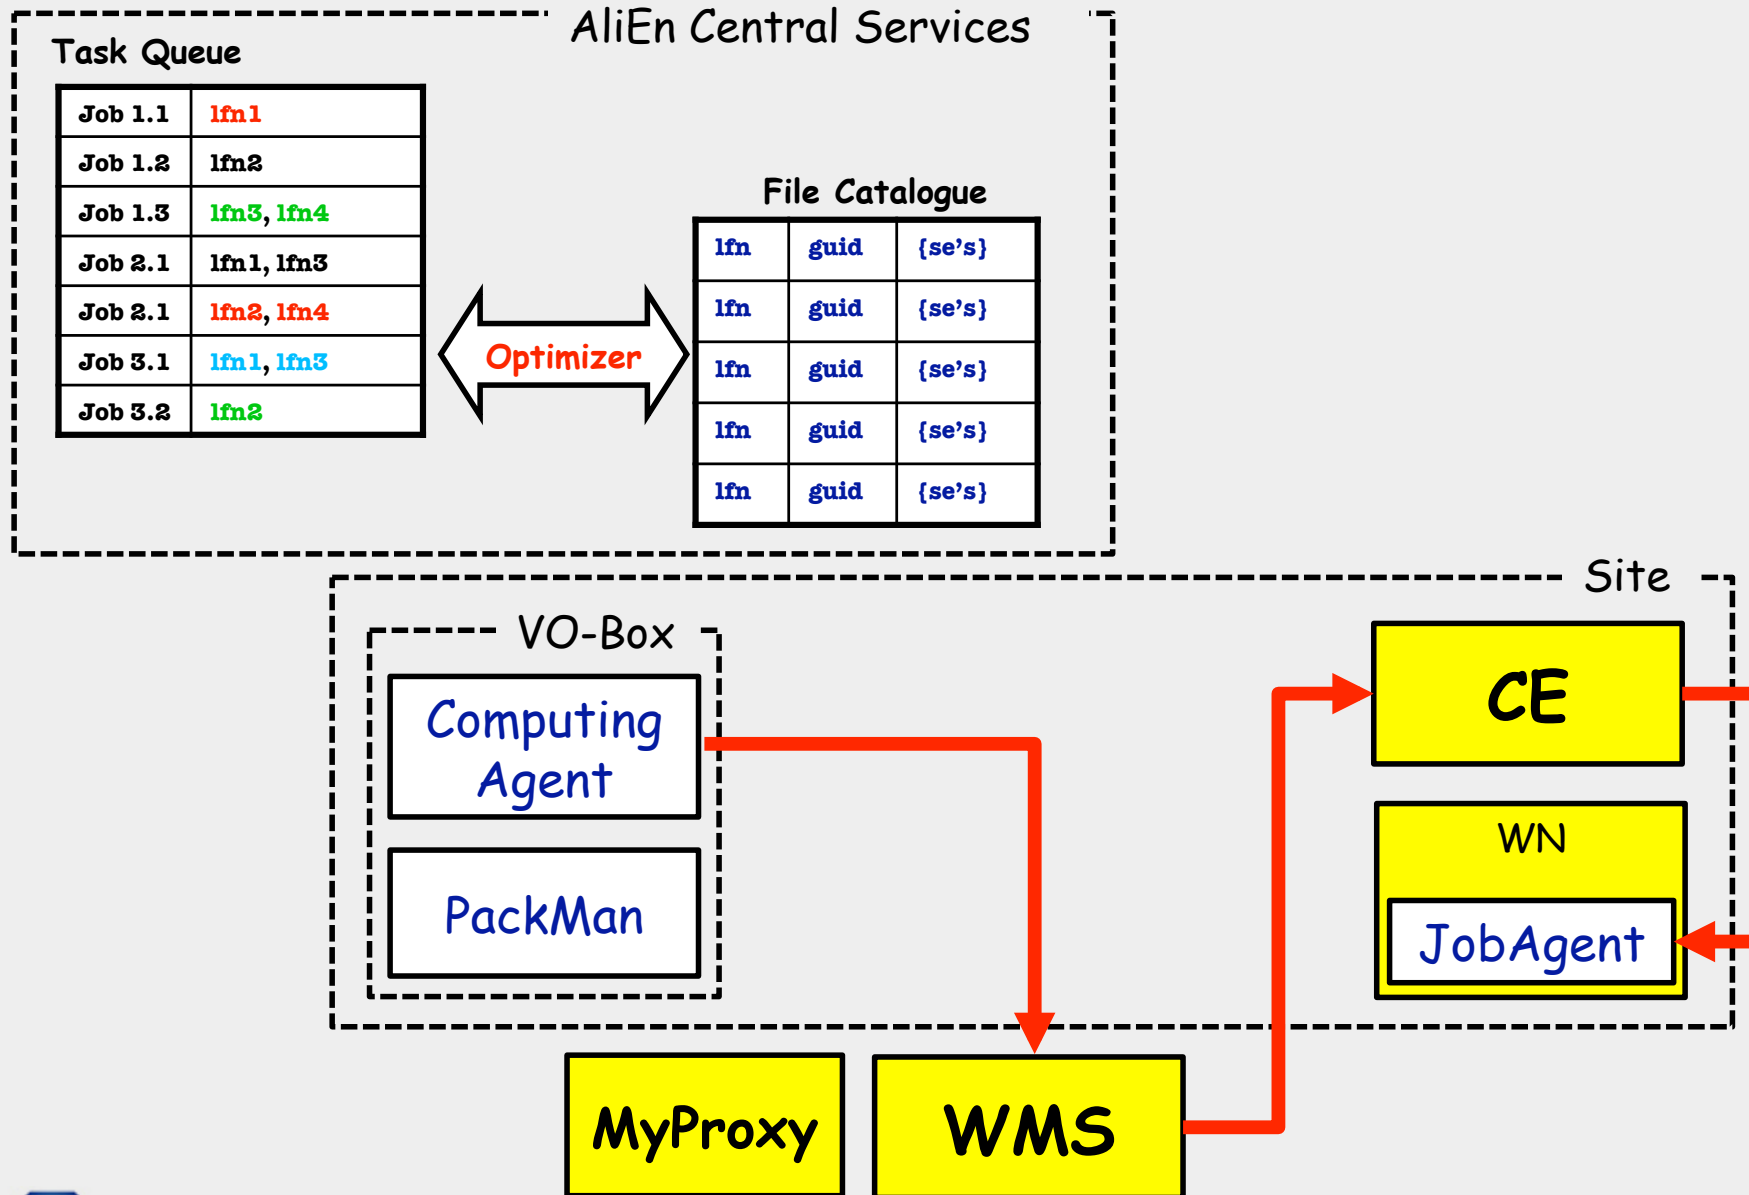


ALICE JOB MANAGEMENT V

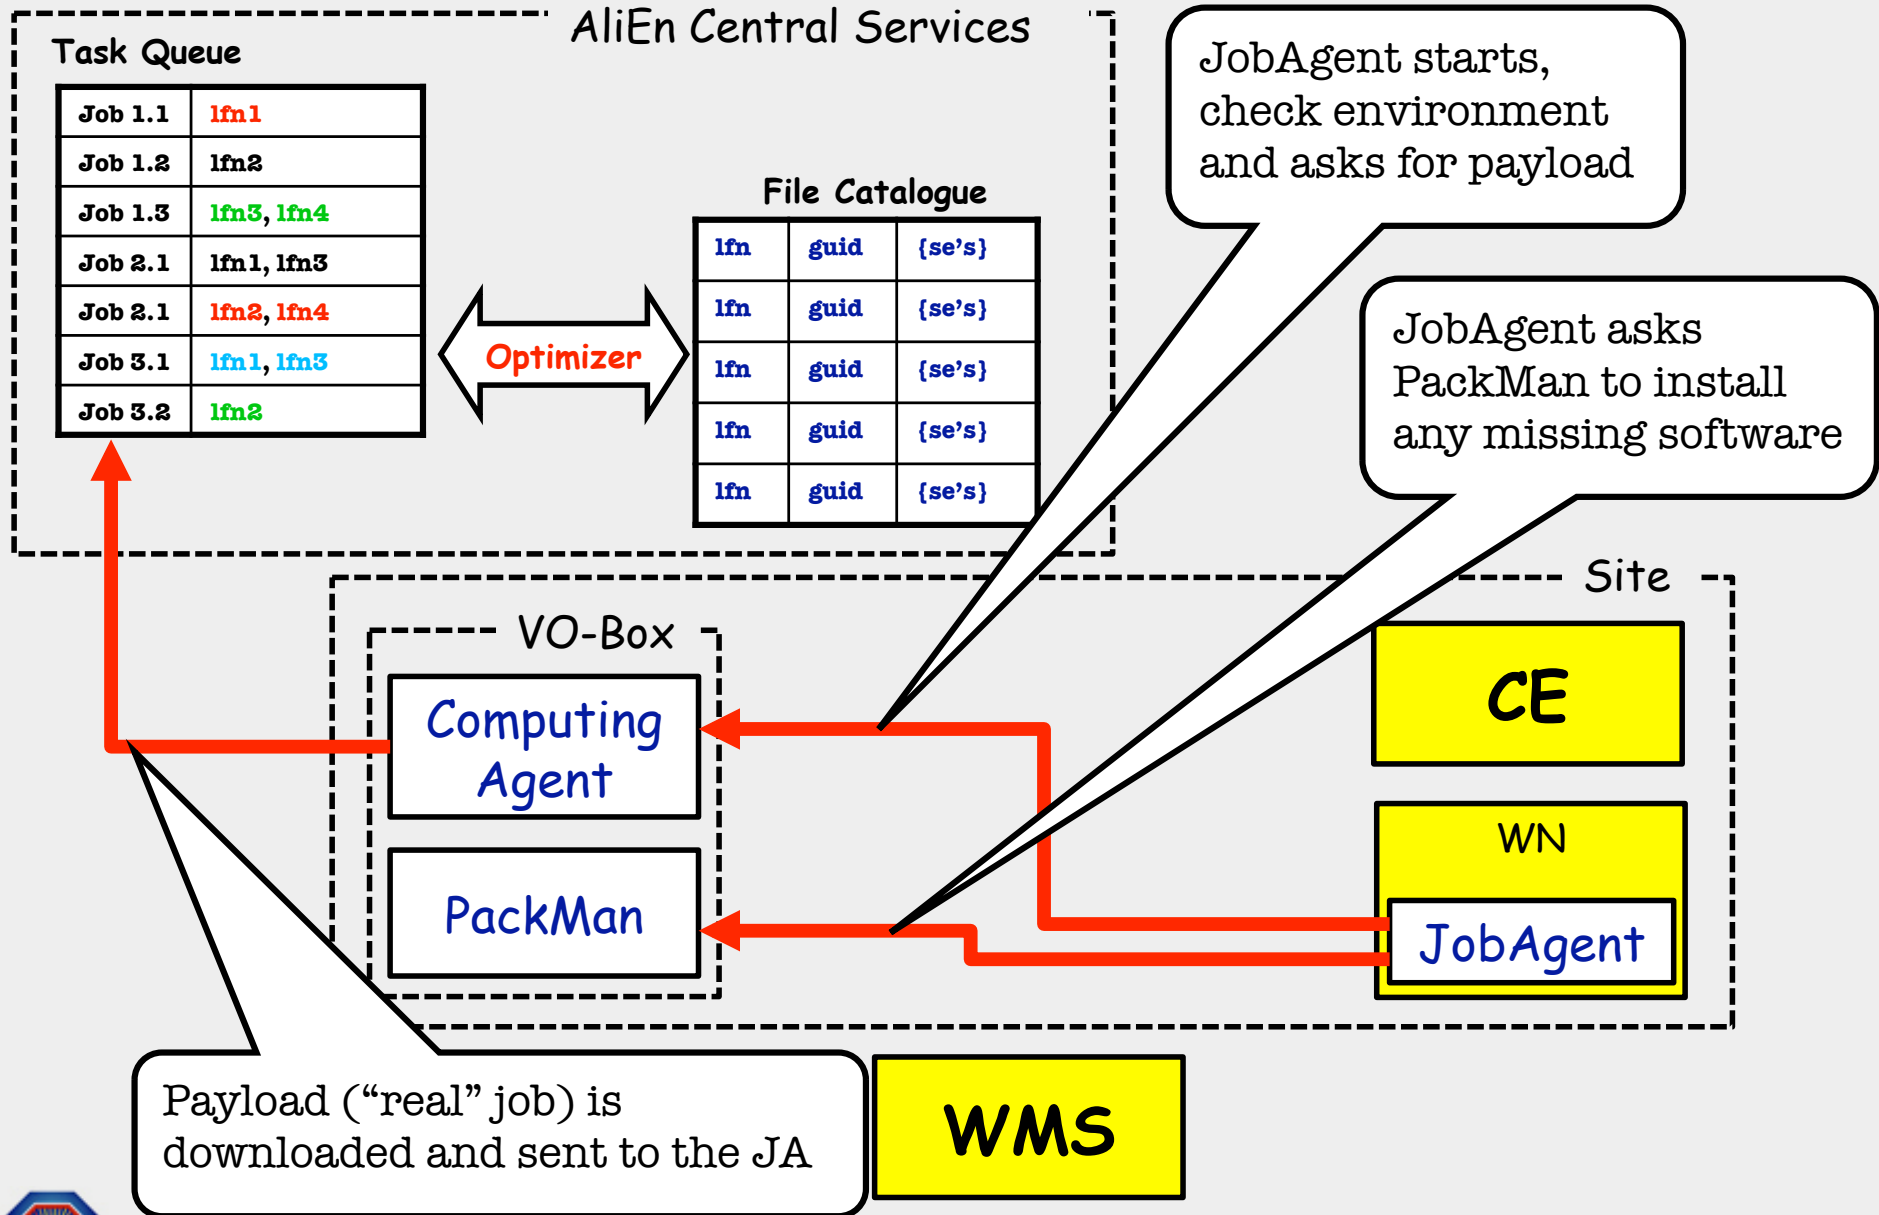


- Task Queue and Optimizers
 - Central DB of jobs to be executed
 - Optimizers split and arrange jobs according to input data, priority policies and/or user defined criteria
- Site VO-Box
 - Thin interfaces to underlying Grid site services
 - Submits JobAgents to site
 - Takes care of proxy management
- JobAgent
 - Runs on WNs, downloads payload from the TQ and executes it
 - JAs create a “virtual grid” on top of existing Grid infrastructures

ALICE JOB MANAGEMENT VI



ALICE JOB MANAGEMENT VII



JobAgent starts, check environment and asks for payload

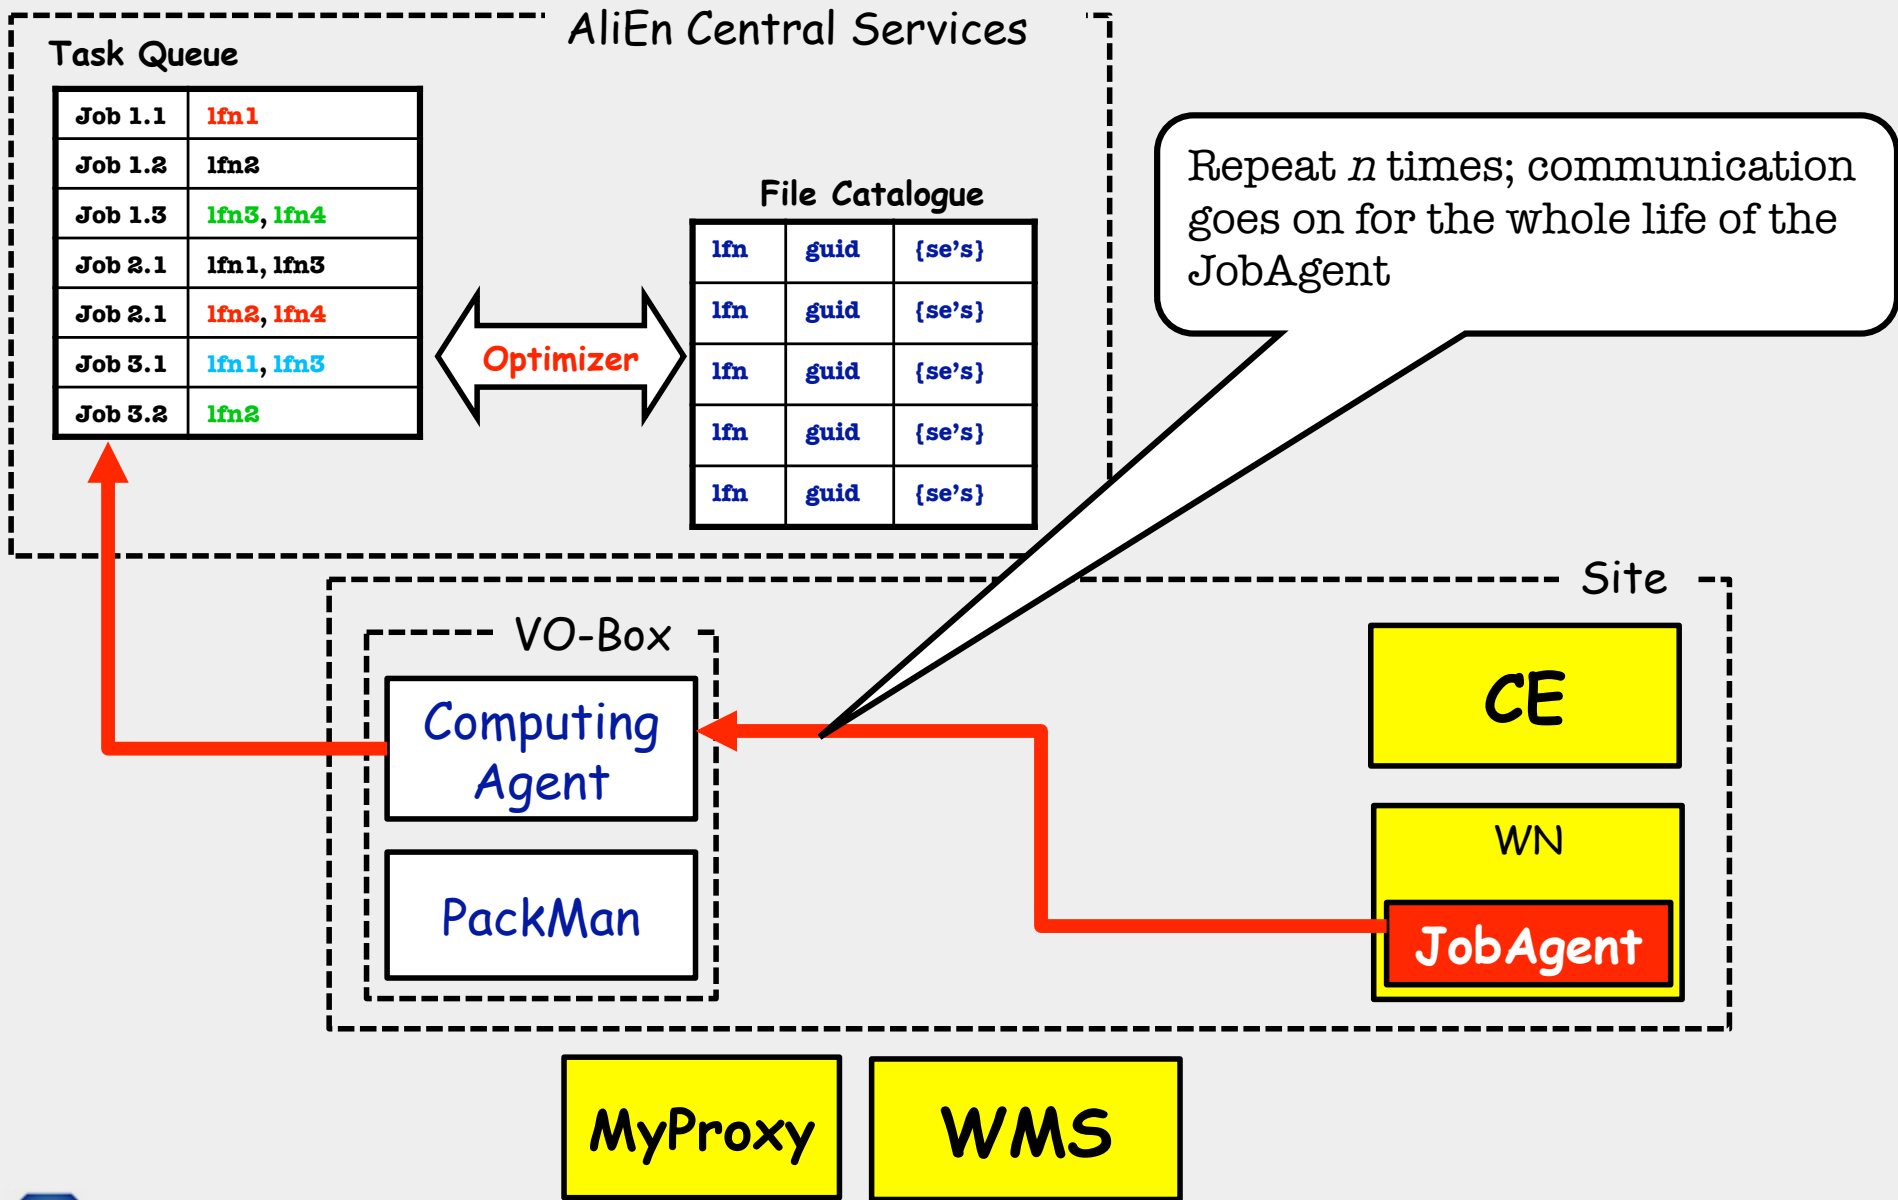
JobAgent asks PackMan to install any missing software

Payload ("real" job) is downloaded and sent to the JA

WMS



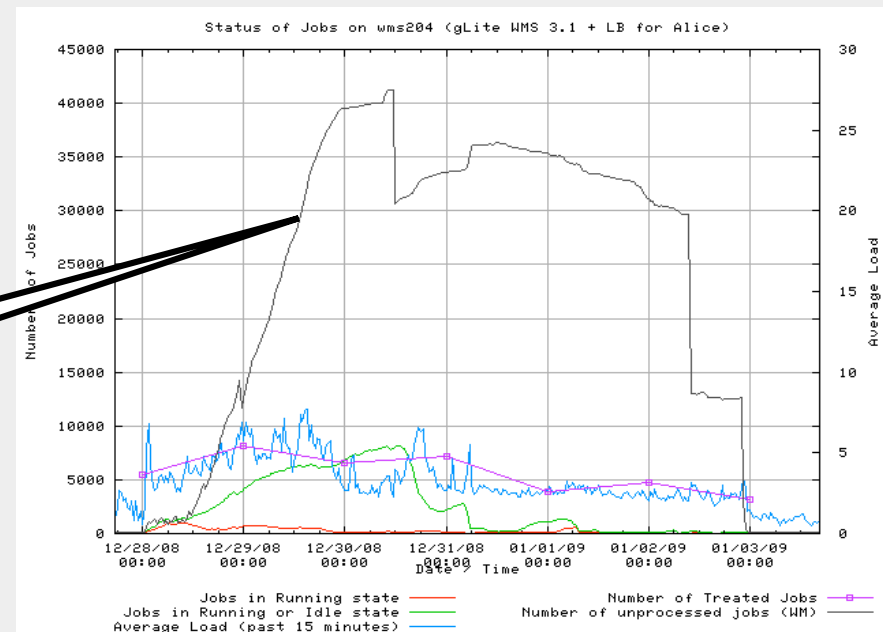
ALICE JOB MANAGEMENT VIII



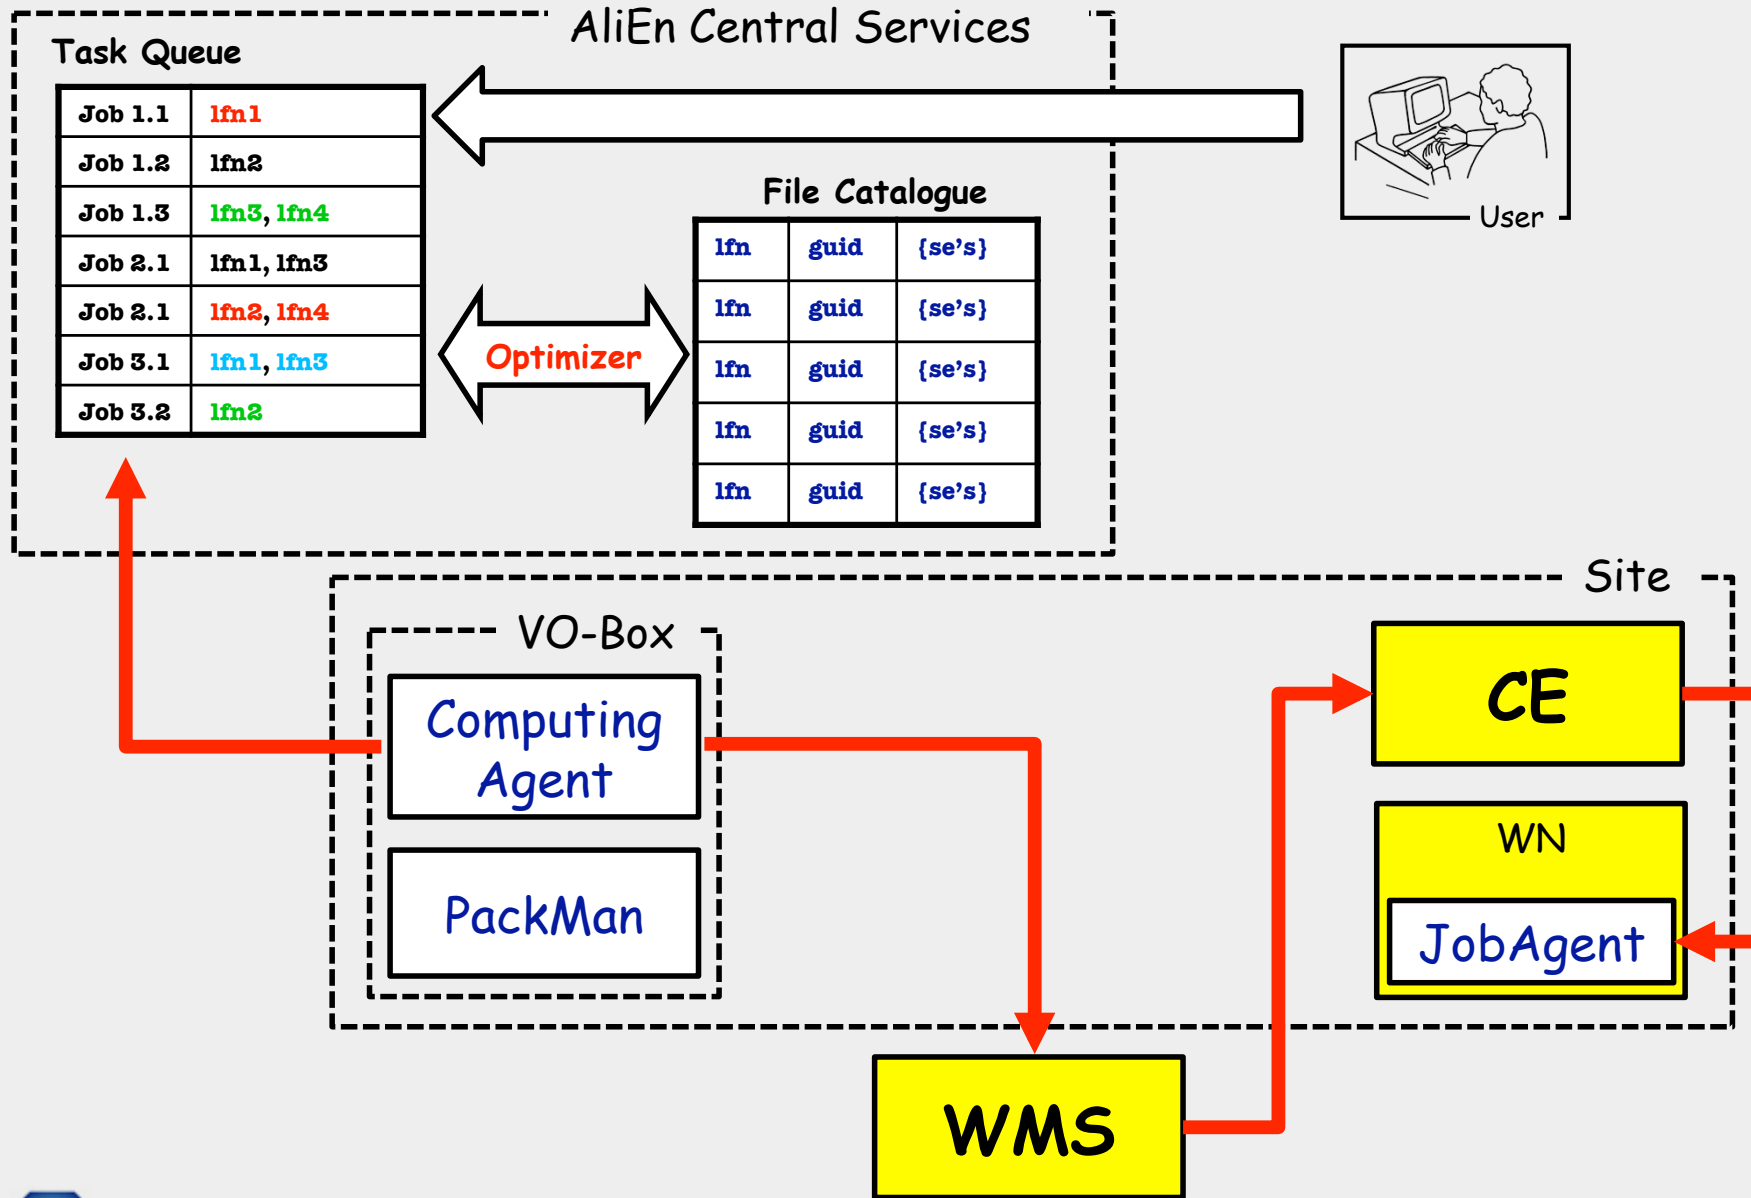
EXPERIENCE WITH WMS

- **Problem:** huge backlogs in the internal queue
 - Reason still unclear
 - Several conditions can trigger the problem
 - These jobs are “invisible” from the IS
- No direct access to WMS internal status
 - This problem is impossible to catch from the VO-Box
 - Only through monitoring tools and site manager intervention

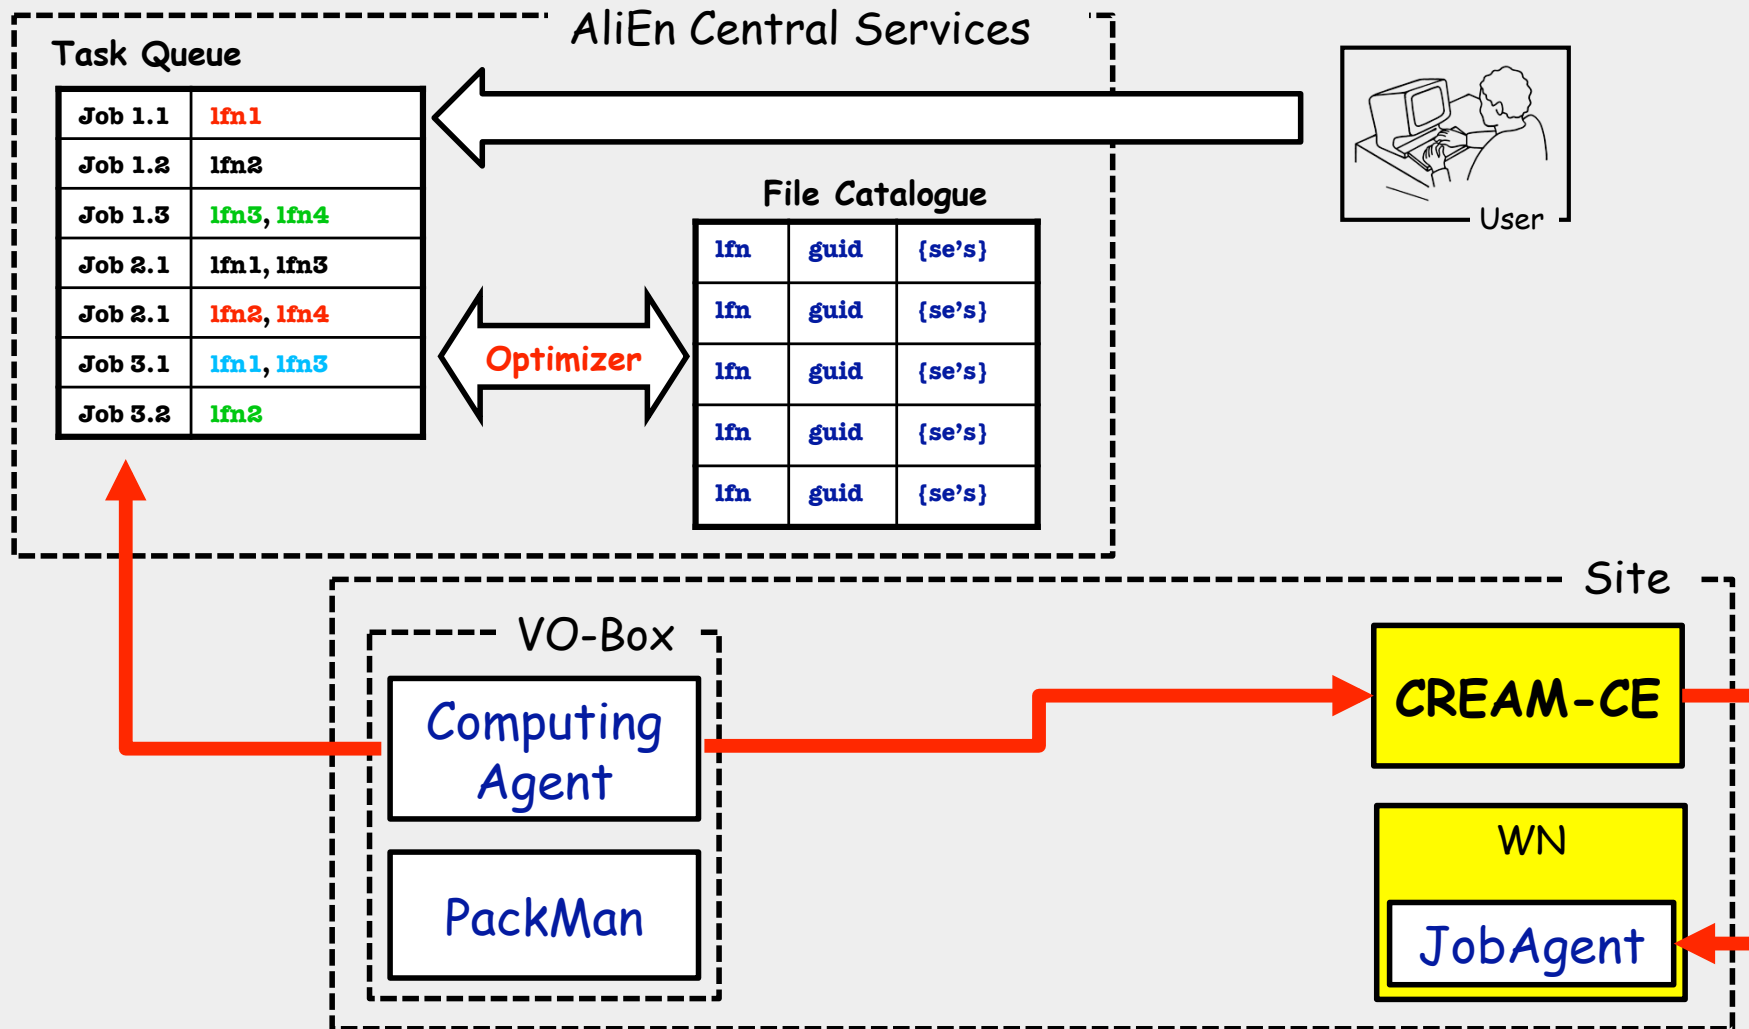
40k jobs in “waiting” status



SUBMISSION THROUGH WMS



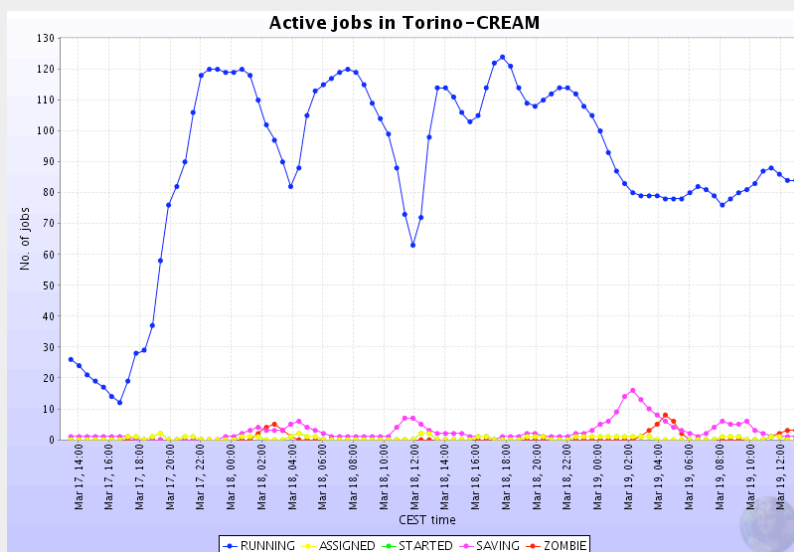
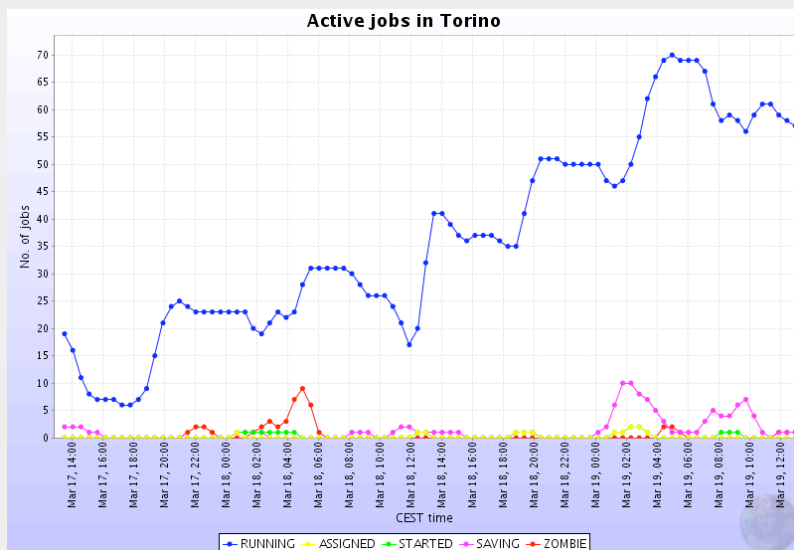
SUBMISSION THROUGH CREAM



- CREAM CE allows direct submission
 - Would have been possible also with older CEs, maybe using Condor
 - Only extra service needed is a GridFTP server on the VO-Box
 - Second VO-Box needed for parallel submission
- First tests in FZK (summer 2008)
 - Very good results
 - Decision to push towards CREAM
- Test phase with 5 sites
 - 67 kJobs in 1 month

- CE installation not always straightforward
 - Many sites reported problems and inconsistencies in documentation
- Excellent support from CREAM developers
 - Thanks to Massimo Sgaravatto
- Afterwards CREAM CE is very stable
- Missing feature: multi-CE submission
 - Simple site-wide load balance

CREAM VS. WMS SUBMISSION



- CREAM and “normal” CEs competing for resources
 - Same WNs, same PBS server physically hosted on the old CE
 - Two WMSs in load-balancing configuration
 - Separate VO-Boxes (similar machines)
- Lower latency translates into steeper ramps
- Factor 2 à regime
 - Please note different vertical scale



CONCLUSIONS

- The overall ALICE job management model is mature and proved itself quite stable and scalable
- The VO-Box + JobAgent mechanism provides a good user and central services insulation layer
- Experience with the gLite 3.1 WMS was disappointing
- Direct submission to the CREAM CE fits well into the schema, and is currently proving itself reliable and very efficient
- Questions, comments?

