

Put some witty  
picture here

## **ALICE Tier1 - Tier2 Workshop**

# **OPERATIONS IN ITALY**

KIT, January 24-27, 2012



# RESOURCES AVAILABLE FOR ALICE

- Tier-1 at CNAF, Bologna
  - Shared with other LHC experiments and a large number of others
- 4 “official” Tier-2 centres
  - “Official” means directly funded by INFN National Scientific Commission according to plans and official pledges
  - Torino, Catania, Bari and Padova/LNL
  - The last two shared with CMS, were officially approved only one year ago
- Cagliari, Bologna, Trieste
  - Local resources, different funding
- CyberSar (CA) and TriGrid (CT)
  - Extra resources for Cagliari and Catania

- 8700 cores, ALICE share is 18 kHEPSpec (about 1600 job slots)
  - Fair share works: last month we used 200% of our share
  - LSF for mgmt
  - Some WNs virtualized in a cloud-like architecture (WNoDeS)
- 930 TB of disk (T1D0+TOD1), 2PB of tapes for ALICE
  - GPFS + TSM for management
  - Xrootd as a front-end protocol
- Staff of 20 (18 FTE) to babysit the whole centre
  - Plus 1 person dedicated to ALICE-specific operations (F. Noferini), not full time
- Also hosts the INFNGrid/IGI Grid Operations Centre
  - More later

# THE OFFICIAL TIER-2S

- Torino
  - Contact: yrs. Truly (ALICE), Stefano Lusso
  - About 1000 cores, 3-400 TB
  - Staff of 3 (not full time, about 2 FTE) + 1 ALICE-specific (not full time)
- Catania
  - Contact: Giuseppe Platania, Andrea Caponnetto (ALICE)
  - About 500 cores, yy TB
  - Staff of 1 + 1 ALICE-specific (not full time)
- Bari
  - Contact: Antonio Franco, Nico Di Bari
  - About 700 cores + CMS, 430 TB
  - Staff of 2 + 1 ALICE-specific (not full time)
- Legnaro (better, Padova/LNL)
  - Contact: Andrea Dainese (ALICE), Massimo Biasotto
  - About 400 cores + CMS, 300 TB
  - Staff of 4 + 1 ALICE-specific (not full time)

(Just joking...)

- Trieste
  - Contact: Stefano Piano
- Cagliari
  - Also CyberSar
  - Contact: Daniele Mura
- Bologna
  - Contact: Francesco Noferini

- The general underlying infrastructure is centrally coordinated
  - INFN GRID “special project” now migrating into the IGI Consortium (Italian Grid Initiative)
    - See next slides
  - ALICE uses a minimum of the common services, so interaction is limited
  - Has several advantages...
    - Site admins can get support for middleware issues
    - Middleware releases are (should be) tested and certified
  - ...and a few drawbacks
    - We have to support a number of other VOs (best effort)
    - We are not completely free in our choices

- IGI is the Italian NGI supported by the Italian Ministry of Research and Education (MIUR):
  - develops and operates grid and cloud services for the Italian grid infrastructure
  - provides services to enable sharing and easy access of distributed resources across sites provided by multiple research institutes and application domains
  - aims at granting the sustainability of the Italian grid infrastructure
  - currently is a EU Joint Research Unit (JRU), based on a Memorandum of Understanding signed by the members in December 2007, and formally supported by the Italian Ministry for University and Research (MIUR) and the European Commission

- All the major Italian Research Organizations (INFN, CNR, ENEA, INAF....), Universities and Computer Centers are part of IGI
- At the end of 2010 the Ministry of Research and Education provided the basis for the IGI sustainability with a long term line of budget:
  - goal: to form a new legal entity
  - the statute of the new IGI Consortium is being finalized
- Skilled people currently working at the JRU will be employed by IGI to preserve know how and the continuity of services





## ● Monitoring and Core Services

- management of central Grid and Cloud services;
- proactive monitoring and problem tracking;
- operations of a middleware certification testbed to test, validate new middleware and foster innovation;
- development of operational tools (IWMSMonitor, HLRMonitor);
- development of a web portal to ease to access to the grid/cloud infrastructure;
- monitoring and accounting services

## ● Helpdesk support

- support to user communities and Grid site managers for middleware installation, configuration and management;
- certification of new sites

## ● Security

- address the operational security-related risks;
- covers all aspects of operational security including Security Incident Coordination and Security Vulnerability Handling in collaboration with EGI and the GARR-CERT

## ● Network

- development and deployment of a set of high quality monitoring tools in order to achieve effective troubleshooting of network problems;
- coordination of the network support;
- interface between European Grid Infrastructure and network providers

# MANAGEMENT AND COORDINATION – 1

- Tier-1 monthly Management Board meeting
  - Includes Tier-1 services representatives (farming, storage, operations etc.) and experiment delegates
  - “Run coordination” report with resource usage and actions status
    - «Why is ALICE not using any tapes at all? What? Still 20% CPU efficiency? Aren't you doing something about this? Why?»
  - Report about resource procurement tenders
  - Experiments reports about current activities and plans for next month

## Tier-2 (and Tier-1) coordination

- Modeled after the overall operations schema
- Alice-it-t2@lists.to.infn.it dedicated mailing list
  - Site admin support
  - General announcements
- (More or less) regular phone conference
  - Right after the TF phone conf
  - Started weekly, now bi-monthly

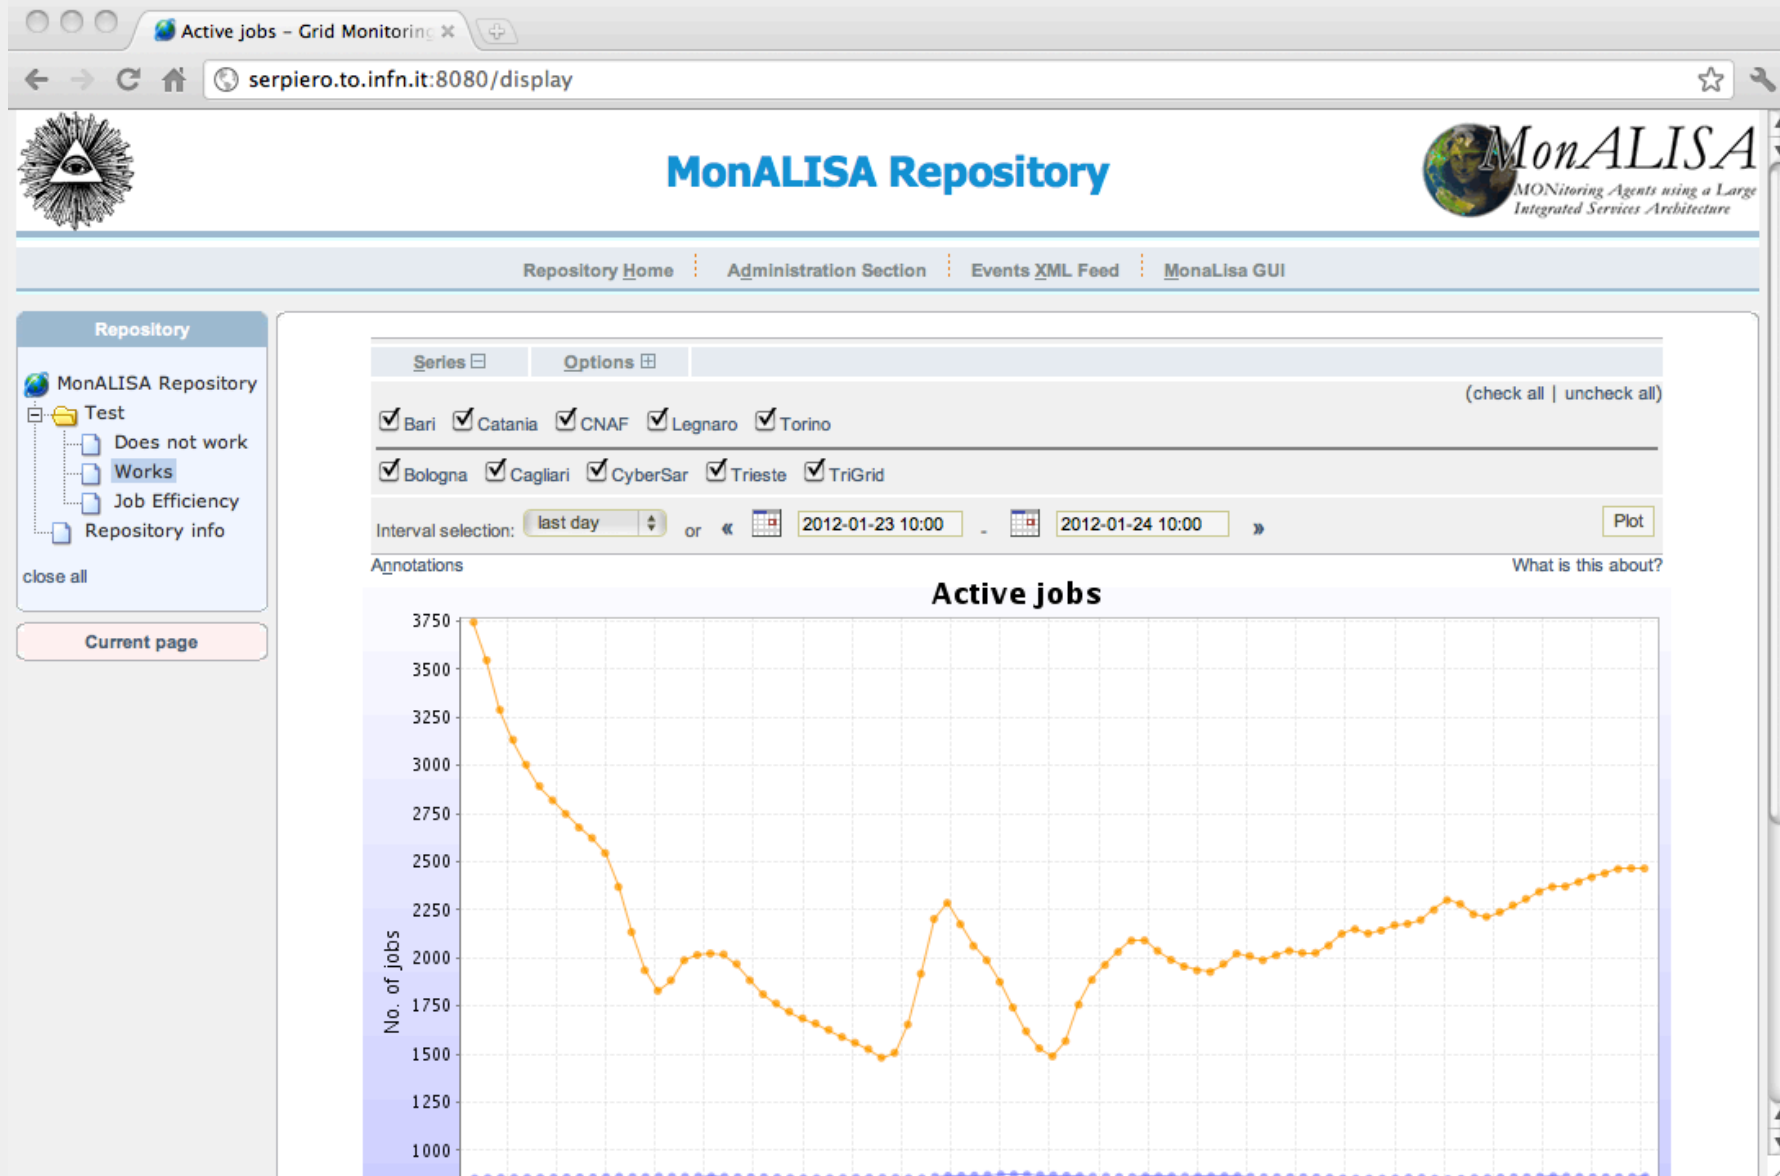
# CENTRALIZED PROCUREMENT

- **Until 2010:** free-for-all
  - Each Tier-2 does its own tenders
- **2011:** gang together Tier-2s
  - ALICE + CMS, since they share some sites
- **2012:** Tier-1 does the job for everybody
  - WNs (ongoing) and storage (upcoming)

# CENTRALIZED PROCUREMENT

- Some obvious pros and cons:
  - Larger tenders mean lower prices
  - But administrative overheads mean longer procedures
- Some less obvious:
  - In larger tenders you have (in practice) less control on what you buy:
  - Inhomogeneous farms (this year we buy HP, next year we buy Dell)
  - There is a subtle trend to buy cheaper hardware, which means more need for manpower in already understaffed Tier-2s
- Machines for services, network equipment etc. are still bought independently
  - Better be kept this way!

# DEDICATED ML REPOSITORY



# DEDICATED ML REPOSITORY

- More detailed monitoring
  - Would be loved by many site managers
  - ML GUI is not felt to be very usable (maybe I'm wrong)
- Really easy to set up
  - Thanks Costin for the support
  - S. Lusso easily set up a second one for PANDA
- Really difficult to configure
  - I'm not sure who to blame 😊
  - We still did not find the time to set up the pages we want...
- Alternative approach?
  - API for download of selected data
  - Obviously cannot be free-for-all
  - Can something like this be provided?



- VO-Box is quite stable
  - Auto-restart usually works
  - Only hitches are with multi-CE, multi-farm configurations
  - Difficult to simulate error conditions
  - Missing a multi-CE test site
  - Will try to provide one in Torino with a “virtual” second farm
  - Would it be possible to merge again CE.log and CE.log\_err?
- Xrootd is a bit less stable
  - Many sites have Lustre or GPFS underneath xrootd
  - Sites should catch FS errors **before** the experiment does
  - But local/regional know-how to debug nontrivial problems is scarce
  - Still heavily depending on core team

- “Regional” site admin workshops/tutorials proved to be very useful
  - Smaller attendances allow for hands-on approach
  - Last one was in 2008
  - Much less problems now!
- Local “informal” user support
  - More formal use of INFN GRID ticketing system never worked
  - 2 tickets generated for support group “ALICE” in a few years (and one was “Test – please ignore”)
- This year will try different approach
  - Two-day AliEn tutorial for users, including a separate session for site admins (exact details tbd)
  - Both for ALICE and PANDA
  - To be scheduled in late spring or early fall
  - Will see how this works

**THANKS!**

- Questions?

