CNAF Experiment Support

Claudio Grandi (INFN Bologna)

The T1 Computing Centre

- Surface: 1000 m²
- Electrical power: 5 MVA
  - Up to 1.4 MW for IT (600 kW in use)
- CPU: 16K cores (195 kHS06)
  - Managed by LSF
- Disk storage 13 PB
  - GPFS file system
- Tape: 1 library with ~20 PB
  - Up to 85 PB with T10Kd
  - HSM system on top of TSM
- 3x10 Gbps WAN connections
- 21 FTE
- Supporting > 25 scientific collaborations

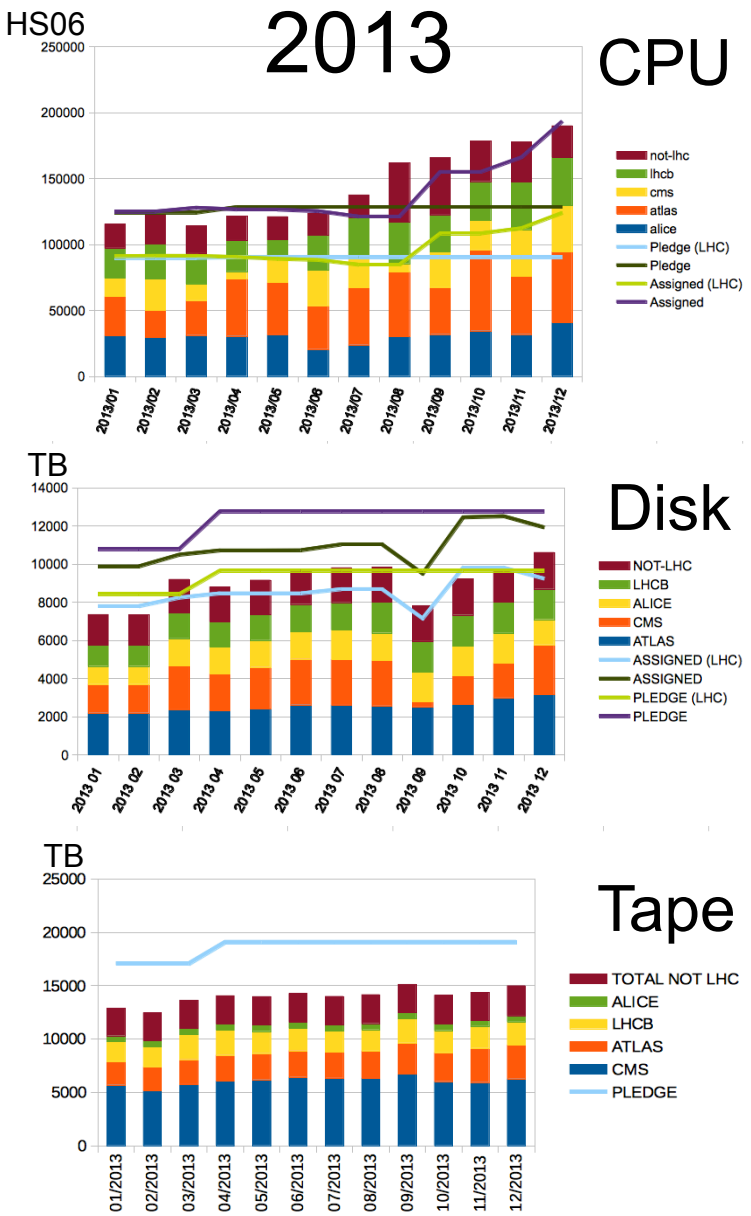

The Tier-1 and Grid Services

- **Tier-1**
  - Infrastructure: manage power, cooling, logistics, security systems, …
  - Farming: operate the bulk of the resources (Worker Nodes, User Interfaces, other services, …) and the user database
  - Storage: operate the disk and tape storage services, and the data transfer services
  - Network: local and wide area networks, IT security
- **Grid services** (ex-IGI)
  - Operations: operate all central Grid services, coordinate middleware deployment in Italy, ticketing systems
  - User Support: Help Desk and support for applications migration to the Grid
  - Middleware: Development, maintenance and support for Grid components

The User Support Service

- Goals:
  - Provide research groups with specific support for day-by-day operations on CNAF resources
  - Support experiments in the definition and evolution of their computing models
  - Foster knowledge transfer from bigger experiments to smaller ones (or simply with less experience in using the CNAF resources)
  - Contribute to CNAF and to supported experiments activities
    - In agreement with national experiment computing coordinators
  - Foster synergies between local and central components of the experiments
  - Train the service's staff (currently 6 *Assegni di Ricerca*)
    - Both for internal needs and for their future career
- Weekly meetings with the Tier-1 Grid Services staff

Ticketing system

- New ticketing system available on https://issues.infn.it/
  - Project: Supporto Utenti T1
  - User friendly interface in development

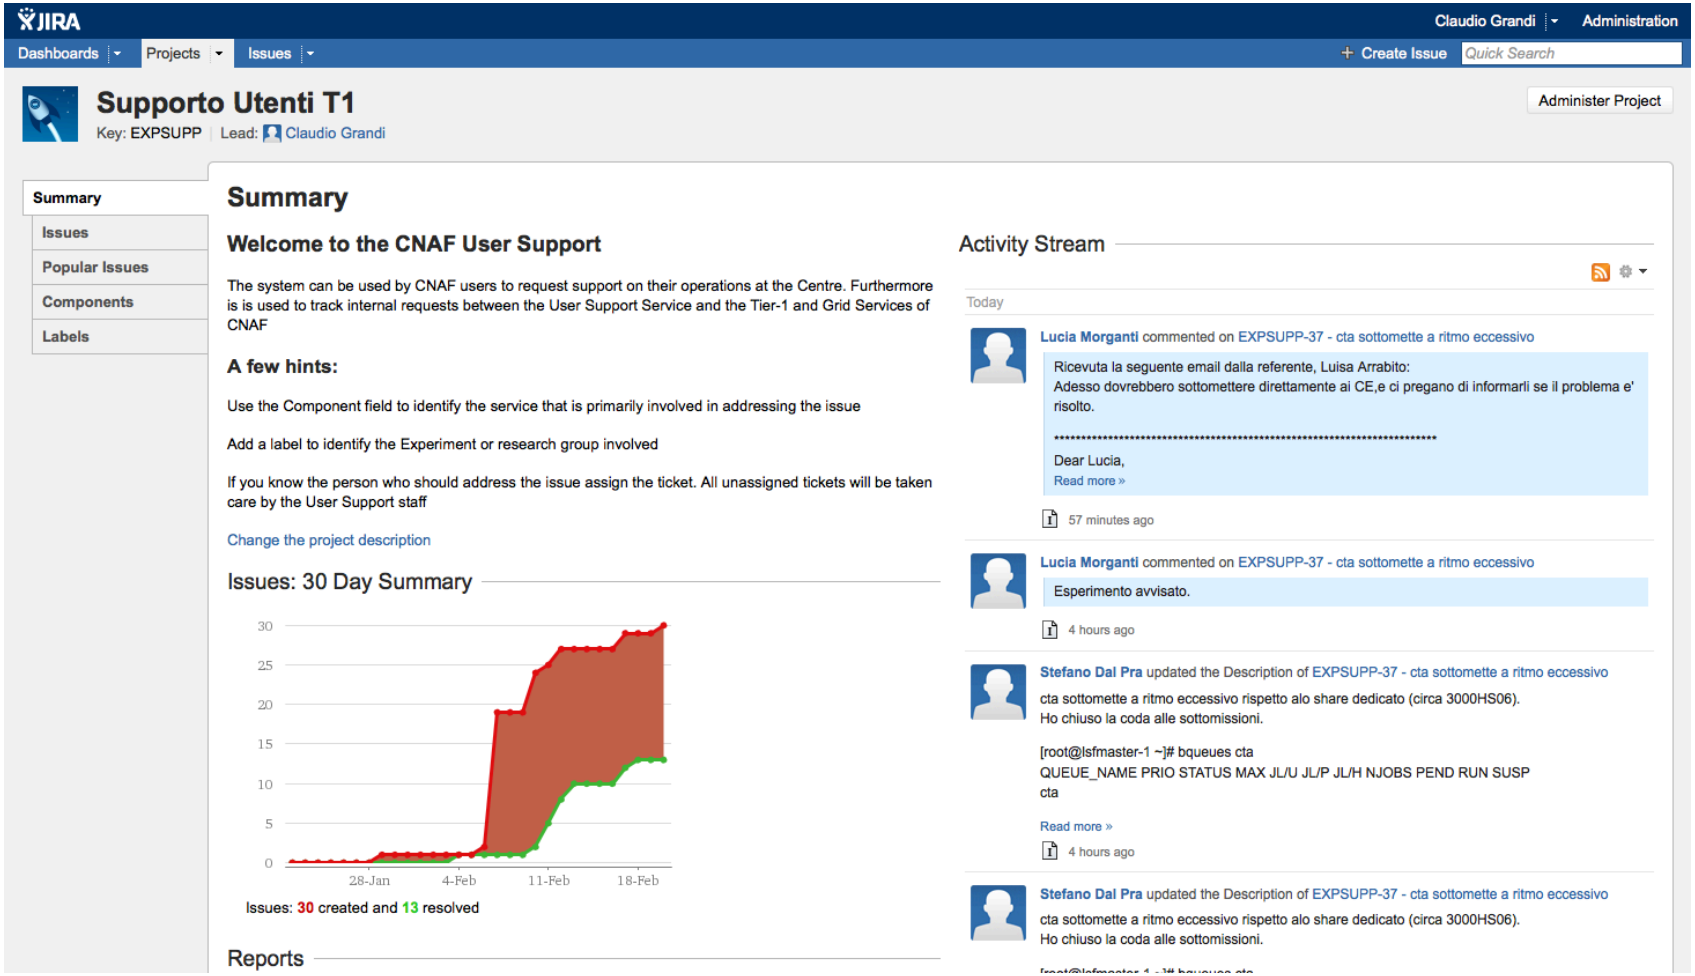

Supported research communities

- The LHC experiments (ALICE, ATLAS, CMS, LHCb) are the main CNAF users
- Most of the non-LHC research groups are astro-particle and neutrino physics experiments
  - CDF, (Babar), (SuperB), Belle II, NA62, KLOE, LHCf, Argo, AMS2, Auger, Icarus, Fermi/GLAST, Pamela, Borexino, Xenon, Gerda, Virgo, CTA, Opera, Magic, Agata, (Panda), (KM3NeT), Theoretical physics groups

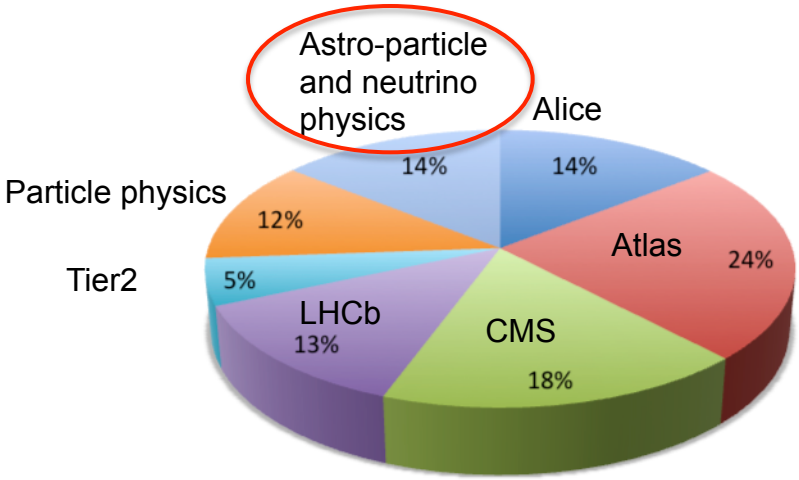

CPU shares

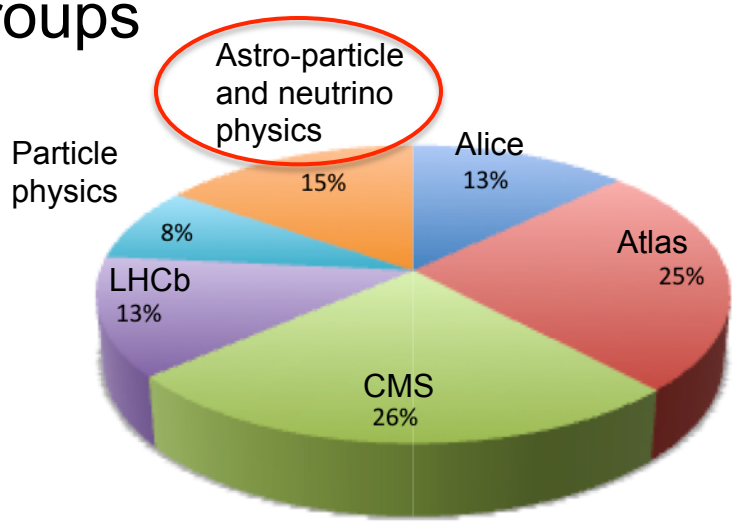

Disk shares

Experiment support for LHC

- LHC experiments provide most of the resources
  - Over 10K cores (11 KHS06), 9.5 PB disk and 16 PB tape
- Rely heavily on Grid technologies, using standard middleware or through tools developed independently
  - Data management (catalogues, agents)
  - Workload management (pilot job frameworks)
  - Resource monitoring tools
- Require also specific services that are managed by the experiment support personnel that are also experiment members
- Use the resources quite constantly but when not using them the increase in available power for the other groups is big compared with their assignments

Experiment support for non-LHC

- Smaller than LHC experiments but many (and increasing)
- Have heterogeneous resource utilization models
  - Via Grid (like LHC experiments)
  - Partial use (e.g. only for data archival)
  - Hybrid: some services via Grid and some locally
  - …
  - Completely local access
- Sometimes not enough computing expertise in the groups to be autonomous
  - Often start with a local approach and migrate to common tools later (Grid)
- Some groups have non constant usage of resources
  - When not using them, they are used by other groups
  - In general they can use more than their quota when they need it

Criticalities 1/2

- Hardware provisioning
  - Bound to tender bureaucracy
- The Tier-1 has been built to be a Grid site
  - Batch job access through CREAM Computing Elements and data access through SRM based Storage Elements
  - Interactive access limited to the minimum on simple User Interfaces without running services
- Many exceptions to this model
  - LHC experiments already have custom services that they manage directly (running in user space)
  - Local job submission to LSF possible upon request
    - Proliferation of user account
  - Local access to the storage from UIs and WNs
    - Requires attention in order to keep performances of GPFS

Criticalities 2/2

- More exceptions to the model
  - Custom Operating Systems (via WNoDeS)
    - Had some problems recently but supported
  - Hybrid access to data (locally and via Grid) and multiple transfer protocols
    - Authorization issues – support issues for multiple servers
  - Remote access to data
    - Dedicated services, network load
  - Custom configuration of nodes (e.g. packages installed on the nodes and not in the experiment software areas)
    - Bookkeeping issues, complexity, fragility
  - Custom services running on nodes (e.g. databases)
    - Ditto, security issues for services out of site's control
  - Root access to nodes
    - Security issues – tolerated only in exceptional cases

Conclusions

- CNAF can provide research groups with enterprise quality services, with 24/7 support
- Resource sharing allows an efficient usage
- A dedicated user support group is available to provide direct support, help addressing specific requests but also evolve the computing practices
- Best quality of service if using standard tools
- Available to address specific (non-standard) requests but need “negotiation” not to jeopardize the efficiency of the infrastructure
- For LNGS specifically, that is an INFN laboratory we need to identify a set of agreed best practices to provide hosted experiments with efficient support for their activities at CNAF

BACKUP SLIDES

Typical scenario

- Data produced outside CNAF, transferred to CNAF using Grid tools to be archived and reprocessed (either locally or via Grid
  - Get Grid credentials
    - X.509 certificates by standard or on-line Certificate Authorities
  - Grid access to CNAF storage to copy, read and delete files or directories, with possible tape archival
    - Upload, download o 3rd party transfers
    - Client software may need to be installed on user machines
  - Use of a catalogue to store information of files (being) transferred, with associated information for file search and transfer validation (e.g. checksum)
    - Standard Grid tools like LFC, Amga or developed by the experiments
  - Interactive access (via ssh) to a CNAF node to submit jobs to the local batch farm
    - Software may be built and installed in a shared area seen by farm nodes
  - Alternatively Grid submission to the CNAF resources via client software installed on dedicated machines (User Interfaces)
    - Preferred even though more complex
    - Transparent access to similar resources at other sites

Example of data transfer and local processing