PRACE iRODS Workshop

Tier 0 – Users needs and requirements

Stefanie Janetzko, FZJ - September 2012

CAPACITIES

e-infrastructure

Outline

- Operating Conditions
  - PRACE Tier-0 systems
  - User Groups
  - Disk Usage
- General Use Cases
- PRACE Tier-0 Experiences and Use Cases
- Further Selected Example Use Cases

PRACE TIER-0 systems


JUGENE/JUQUEEN

HERMIT

CURIE

FERMI

SuperMUC

User Groups

- Regular / Preperatory Access
- Different needs and prerequisites
  - Kind of access
  - One or more execution systems/sites
  - Academic or industrial project
  - Scientific field

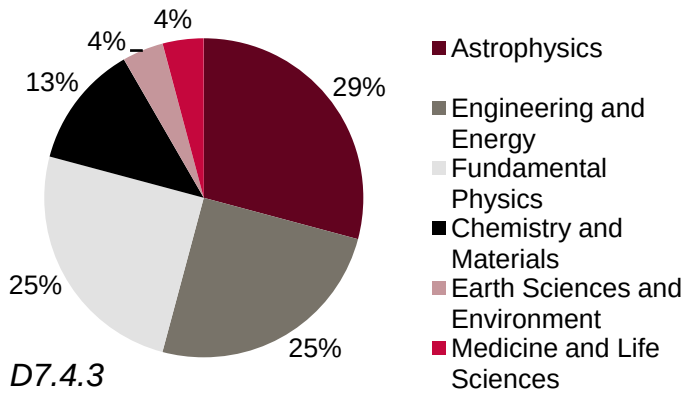

From D7.4.3

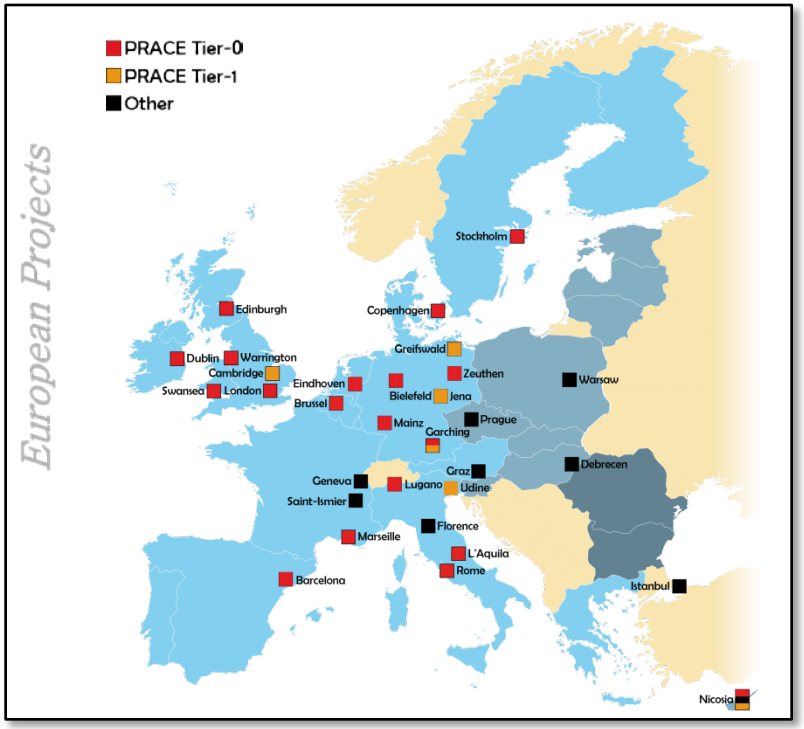

D7.4.3 Tier-0 Applications and System Usage

- Survey on PRACE-Users on CURIE and JUGENE
- Data needs of Tier-0 projects
  - Snapshot Q1-2012: actual usage up to 50TB
  - Wide range of disk usage is observed
  - Small number of projects typically account for most of the disk usage
  - Current applications: up to 100TB
  - Min. disk space per production run:

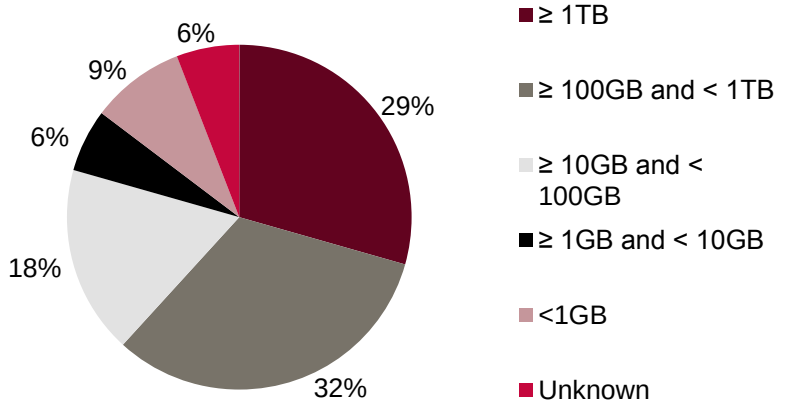

Data Volume – Example FZJ

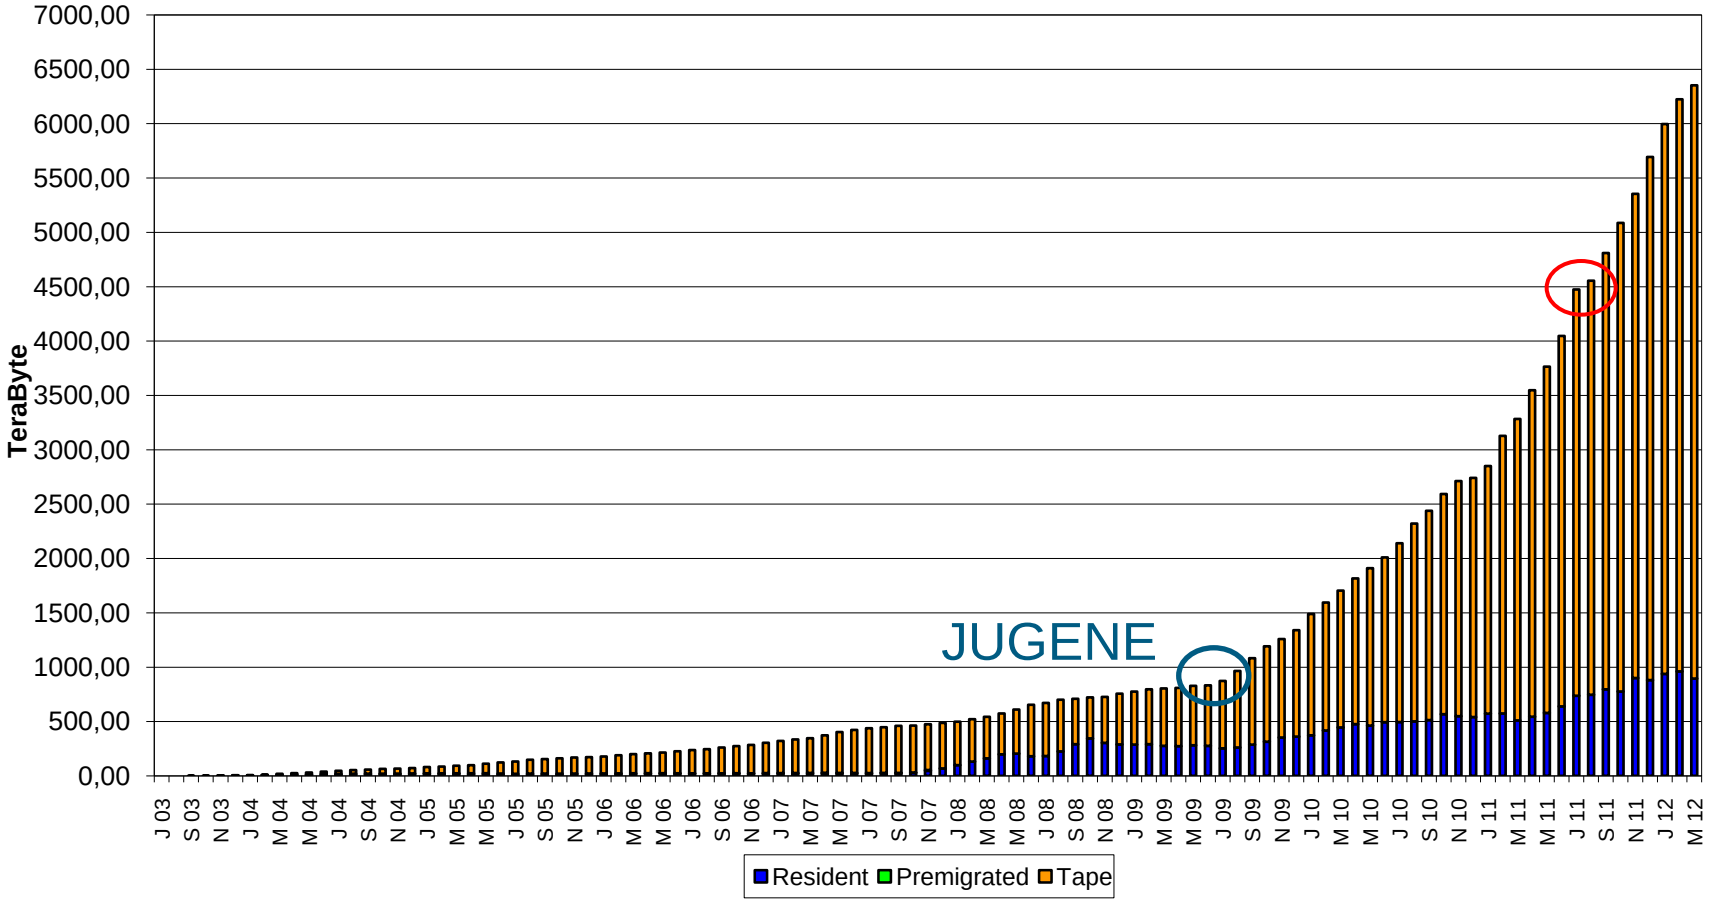

User Requirements

- **Storage**
  - Scratch (no back up, automatic cleaning of old data)
  - Work (no back up, no cleaning, project's lifetime)
  - Home (back up, no cleaning, project's lifetime)
  - Long-term storage (beyond project's lifetime)
- **Transfer**
  - Local / External
- **Accessibility**
  - Single user, multiple user(s), and/or multiple user groups and communities → Authentication, Authorization
- **Reliability**
  - Access and Data reliability




General use cases

Use Case I: Big Simulations

- **Description**
  - Larger HPC systems allow increasing simulation sizes
  - Amount of data produced increases
- **Requirements**
  - Storage: Scratch, Work, Home
  - Transfer of simulation data
- **Challenges**
  - Required storage capacity per project increases on average
  - Enough capacity must be available
  - Connection to compute nodes
  - Bandwidth or long-term storage

Use Case II: Multiple-User Data Access

- **Description**
  - Projects with
    - different users of one community
    - Interdisciplinary projects (users of several communities)
    - different execution systems/sites
- **Requirements**
  - Data transfer to/from the execution systems
  - Storage capacity at the execution systems
  - Accessibility of data for all collaborators from/on all systems
- **Challenges**
  - Increasing amount of data requires increasing bandwidth & storage
  - Multiple user(s)/groups (national, international) → accessibility of the same data on different systems and for different communities

Use Case III: Pre-/Postprocessing

- **Description**
  - Production runs on one system
  - Preparation of runs/Analysis of runs on a second system
- **Requirements**
  - Access to the data from both systems
  - Data transfer between the systems or common filesystems
- **Challenges**
  - Increasing amount of data requires increasing bandwidth & storage
  - Accessibility of the same data on different systems

Use Case IV: Usage of external Data(bases)

- **Description**
  - Projects using data of external databases
    - Meteorology (Climate data)
    - Biophysics (protein databases)
- **Requirements**
  - Connection to external databases
  - (Temporary) storage capacity on simulation system
- **Challenges**
  - Bandwidth to external data sources
  - Storage capacity

Use Case V: Data provision in Standby

- **Description**
  - Long-term provision of data
    a) Storage of data from scientific publications (~10 a)
    b) Amount of data exceeds local capacity available to users
    c) Reopening of projects
- **Requirements**
  - Long-term storage and data reliability
  - Project-independent storage
- **Challenges**
  - Storage capacity
  - Long-term availability and reliability
  - Access (authorization)

Use Case VI: On-demand simulations

- **Description**
  - Fast data access on demand
    - medical purposes
    - emergency cases
- **Requirements**
  - Fast/instant availability of data
  - Confidentiality
- **Challenges**
  - High bandwidth
  - High reliability and confidentiality of data storage
  - Long-term storage
  - Access management

PRACE experiences and use cases

PRACE Tier-0 projects – experiences I

- Main Requirements – regular projects
  - Storage capacity
    - Up to 100 TB
    - Provided by the hosting site
  - Data transfer
    - Transfer needs up to 100 TB (project home <> Tier-0 site)
    - Often mainly at the end of the project, terabytes of data: needs to be transferred continuously → focus already in project application
    - Tier-0 cases: PRACE<>external, PRACE<>PRACE site
    - Tools used: scp, bbcp, (Grid solutions: gridftp, uftp)

PRACE Tier-0 projects – experiences II

- Main Requirements – PA projects
  - More than one execution system
    - But different to Tier-1 case:
      - Test of different systems, no workflow
  - Access to data from different uses
    - Especially type C projects
    - Small amount of data (e.g. source code)

PRACE Use Case Pre-/Postprocessing (UC III)

- Example: different PRACE regular access projects at FZJ
- Special needs of jobs for Pre-/Postprocessing
  - Small, but long jobs
  - Main memory needs
  - Visualization
- Usage of visualization cluster JUVIS and Tier-1 system JUROPA

→No data transfer due to common GPFS filesystem

Further selected example use cases

Example: Atmospheric Infrared Sounder (AIRS)

(Combines several Use Cases)

- NASA provides global climate data
  - About 50 TB data for 10 years measure time
- Use at SimLab Climate (JSC):
  - Continuously retrieving data on JUGENE
    - per day: 2x240 files, 14 GB
    - Currently time-consuming data management
      - Storage on tape
      - Continuous check of download status and data correctness
      - Extraction of data for simulations
      - Data access of other collaborator difficult
      - Data access from different machines desirable

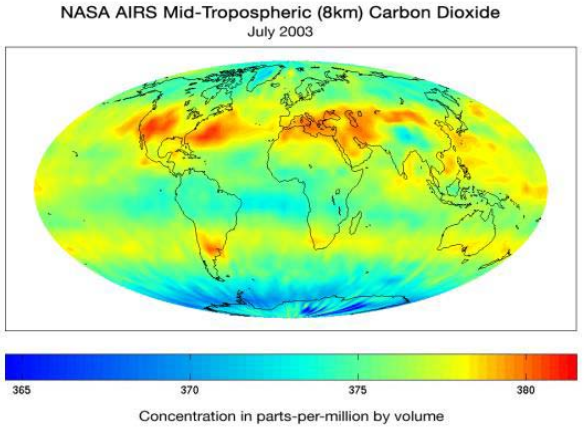

1. Temperature distribution - 1.1 TB input → 20 TB output (hard disk)
2. Daily map of temperature radiation - 0.5 TB output data (hard disk)

Example: Simulating Galaxy Population of dark Energy Universes (Use Case I – Big Simulations)

- Simulation on the JUROPA system
  - Largest simulation of its kind
    - 6.720^3 particles, Cube of 4.1^3 Gpc
  - 12,000 compute cores
  - 30 TB main memory
  - 100 TB of data produces
- Post-processing of data
  - In parts on JUROPA and at RZG
    - ➢ Data had to be transferred to RZG (bbcp)

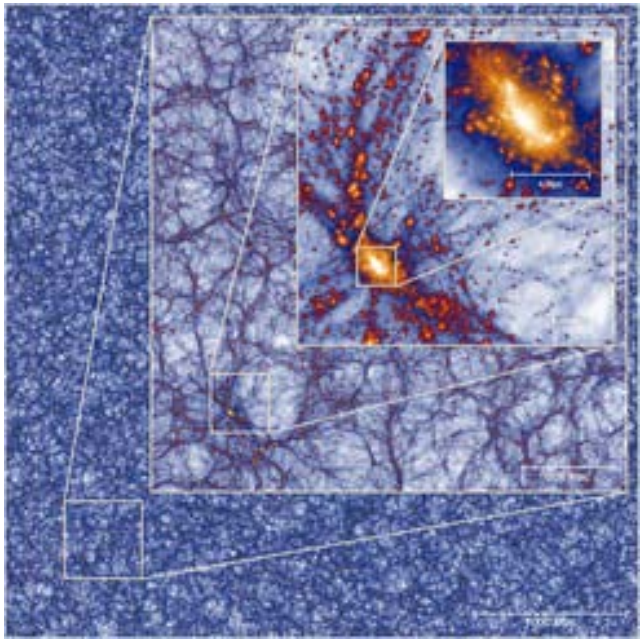

Example: ISS Project AMS (Use Case I – Big Simulations)

- Alpha-Magnet-Spectrometer
- Data analysis on the JUROPA system
  - Input data (transfer from CERN to Jülich)
    - 45 TB
  - Output data
    - 155 TB (~200 TB per year)
- Transfer via Grid-software
  - lcg-utils (LHC Computing Grid data management)
  - First transfers of 15TB at weekend

Thanks for your attention!