Ian Bird

LHCC Referees' meeting;

CERN, 3rd March 2015

Project Status Report

Pledge usage

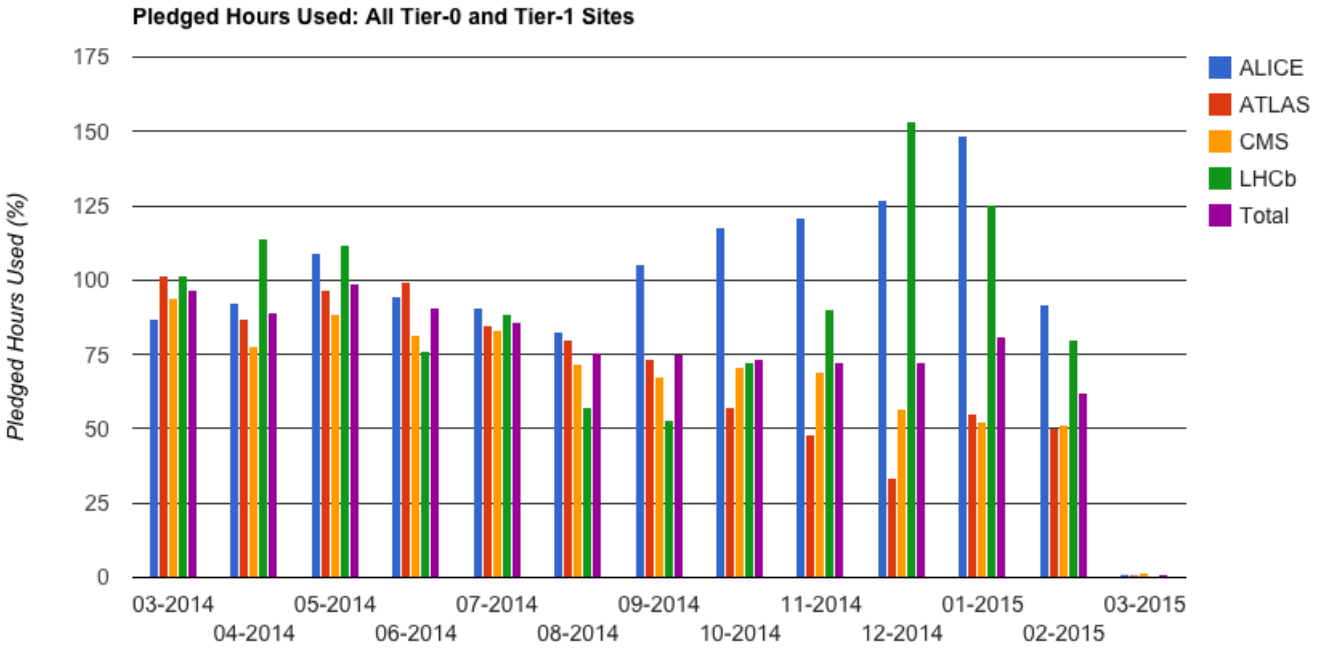

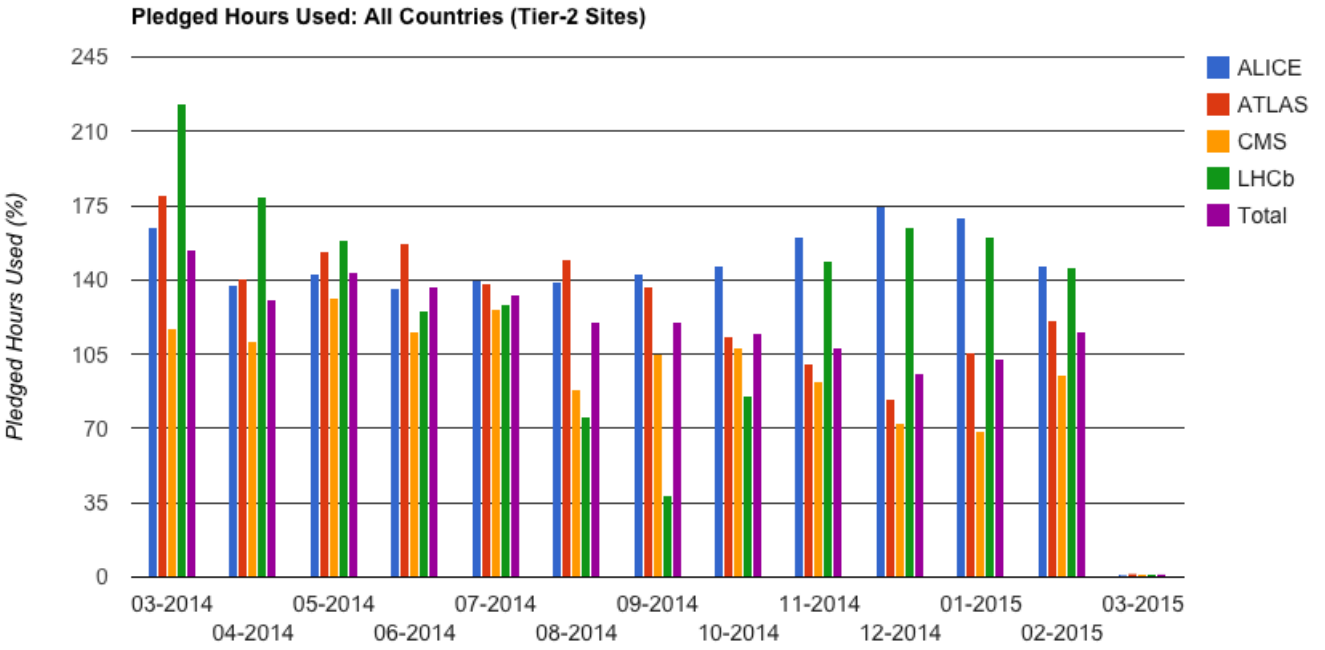

YEARS/ANS CERN

WLCG

New Tier 1s, providing significant contributions Nov 2014 – Feb 2015

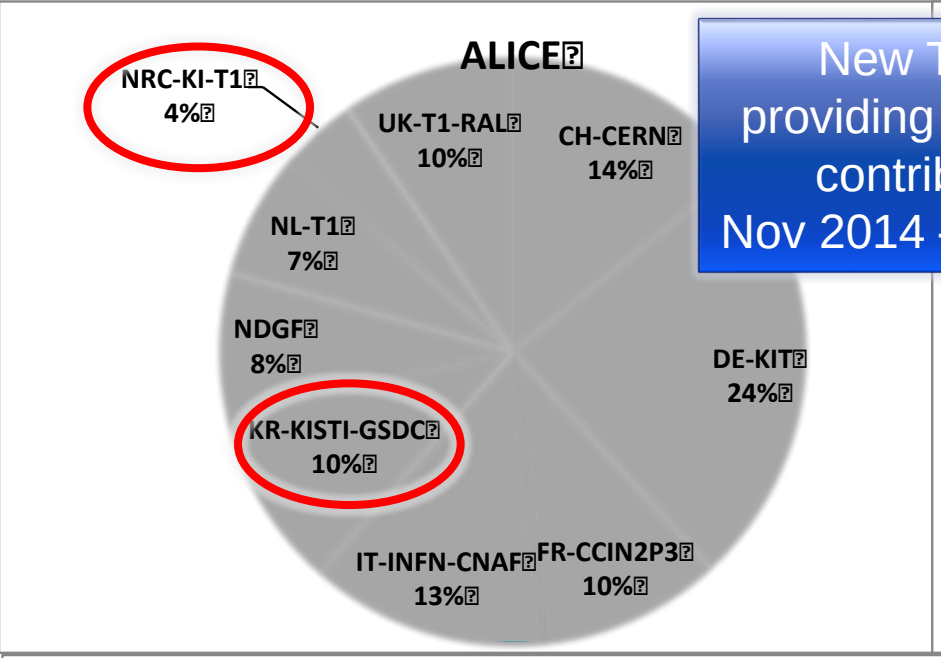

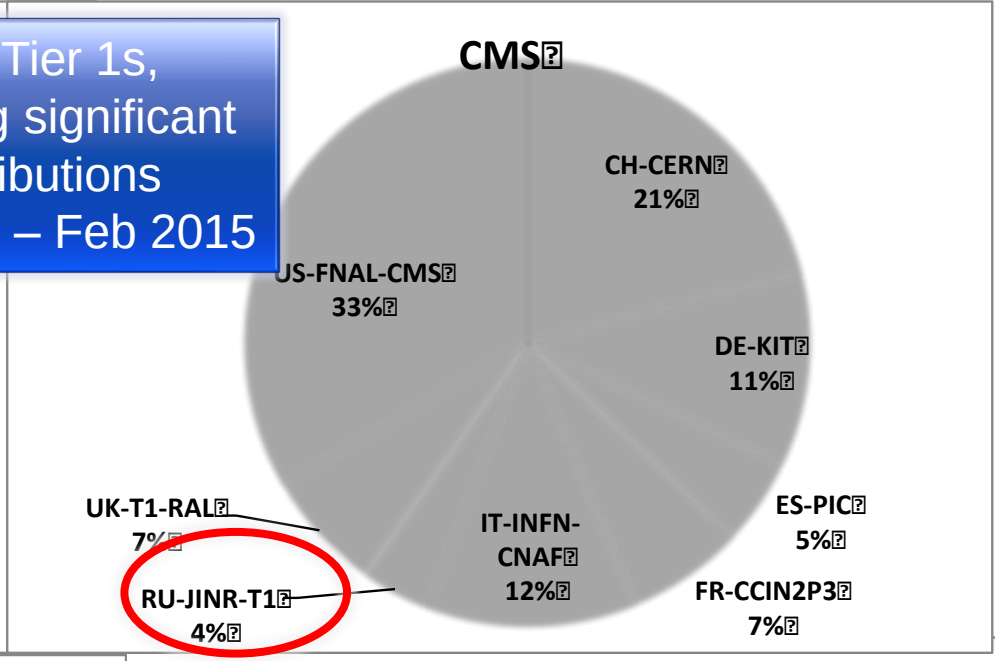

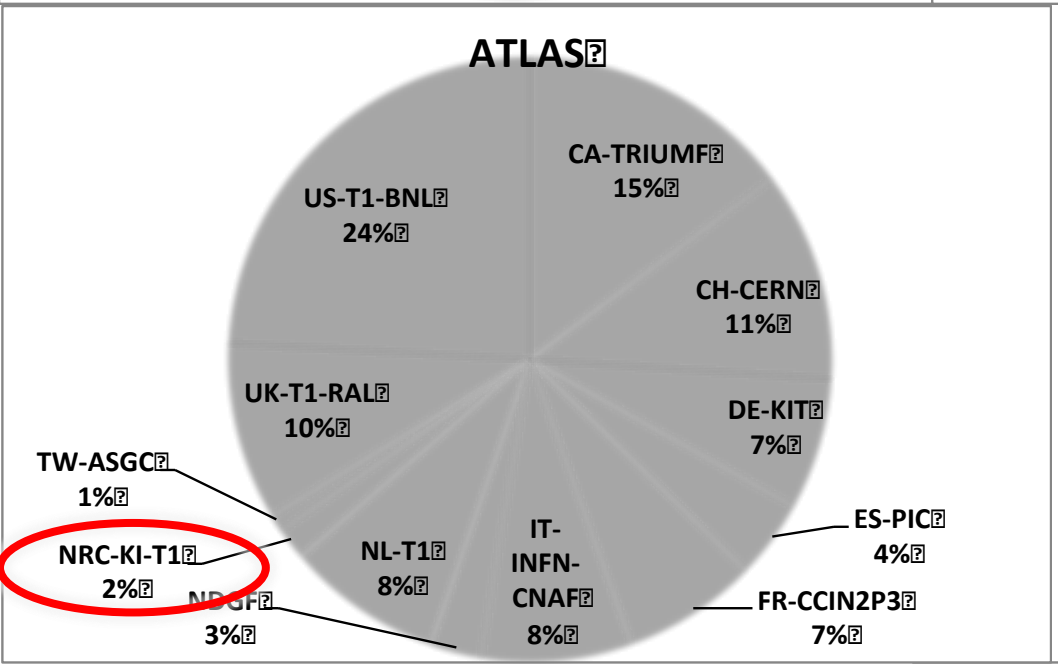

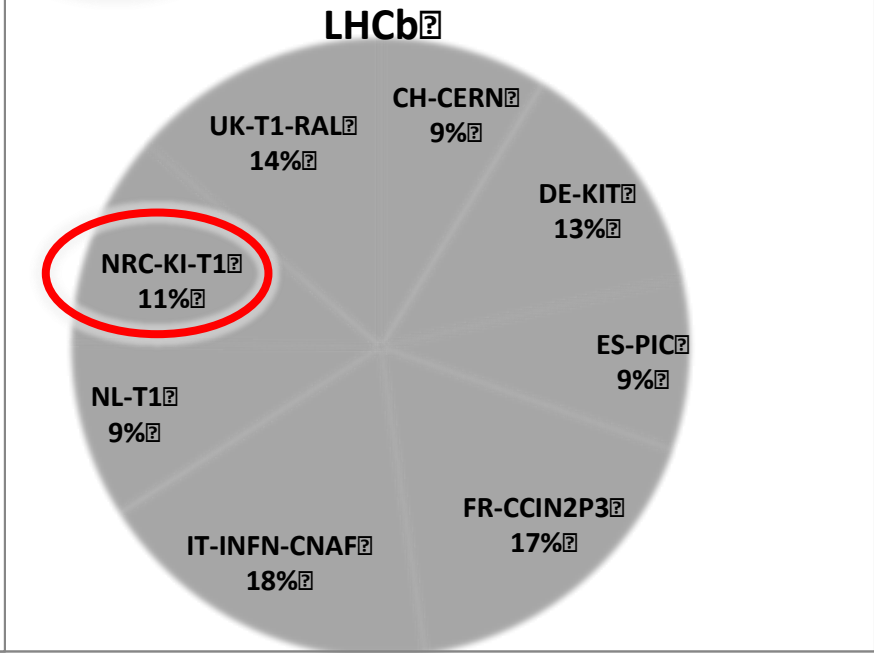

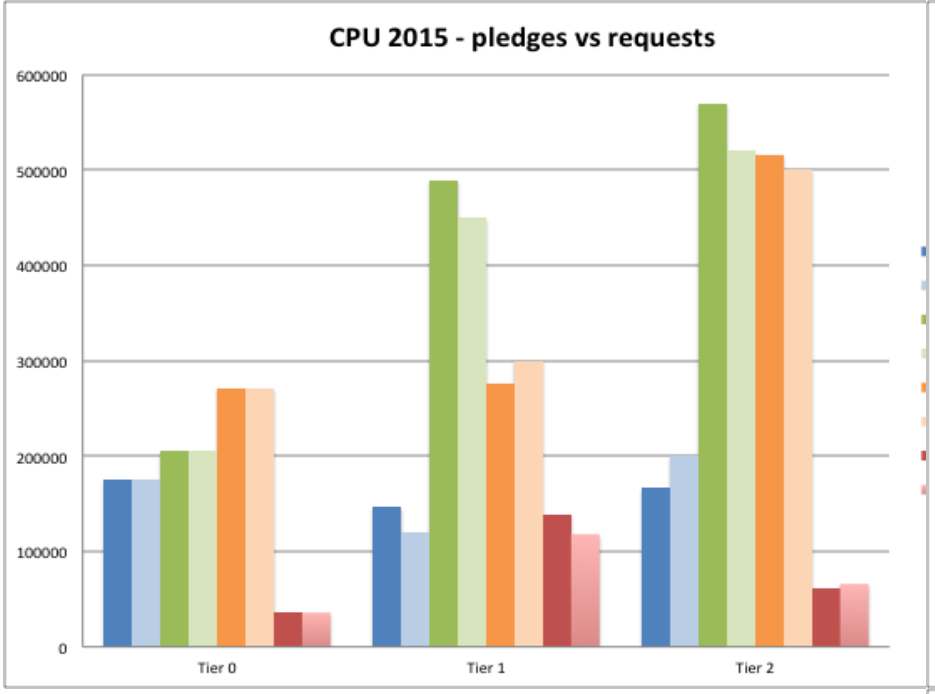

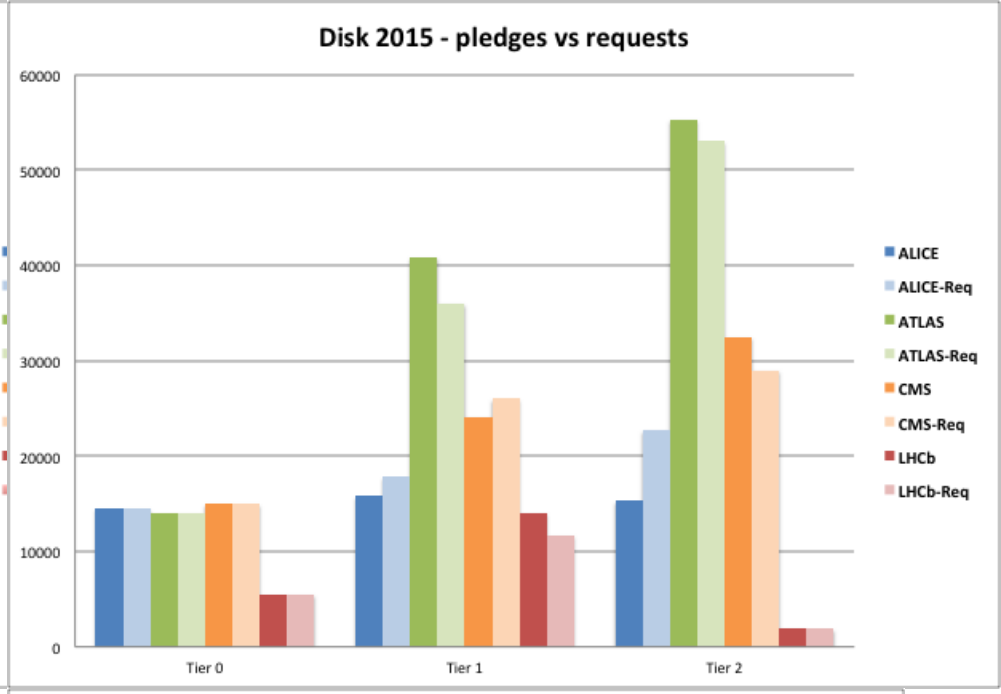

Pledges vs requests for 2015

ALICE
ALICE-Req
ATLAS
ATLAS-Req
CMS
CMS-Req
LHCb
LHCb-Req

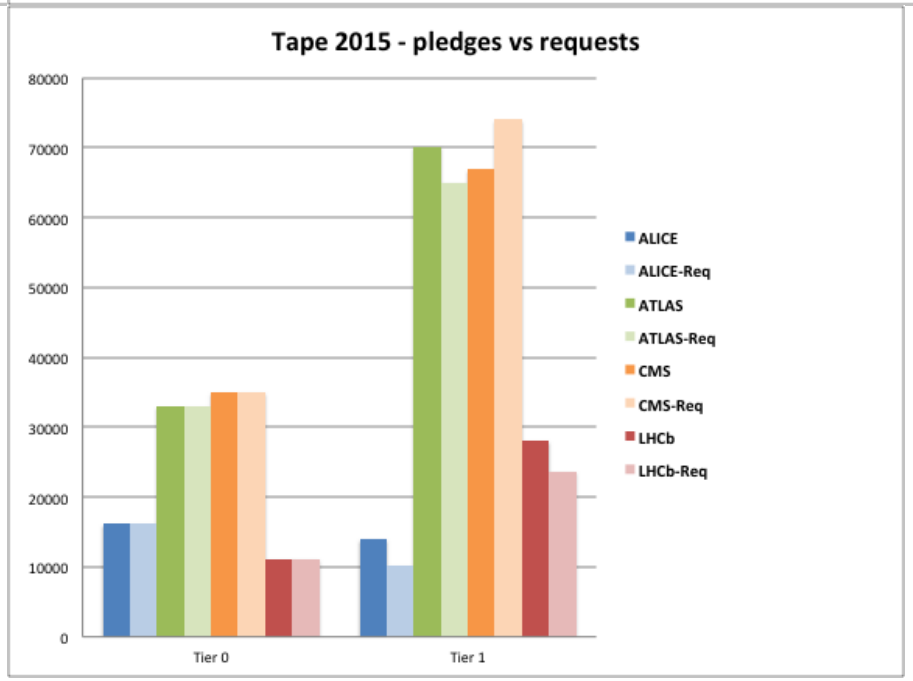

Resource installation

- Mostly on track to be in place for April
- One or two exceptions
  - But expected by ~June
- Not a real problem for 2015, given LHC startup schedule

Tier 0

- 2015 capacity
  - 50 PB disk, 750 kHS06 (36 k cores)
    - 2/3 in Wigner, 1/3 in Meyrin

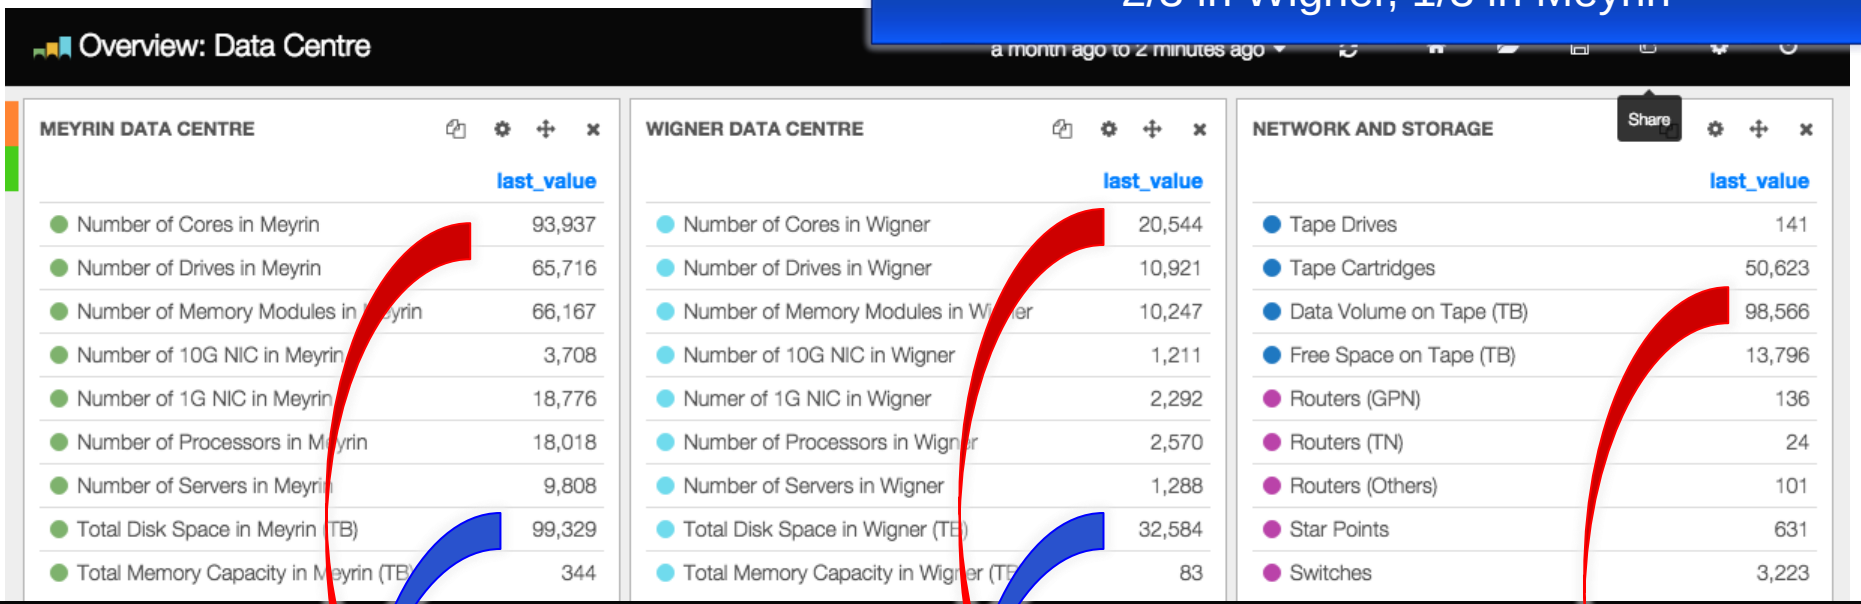

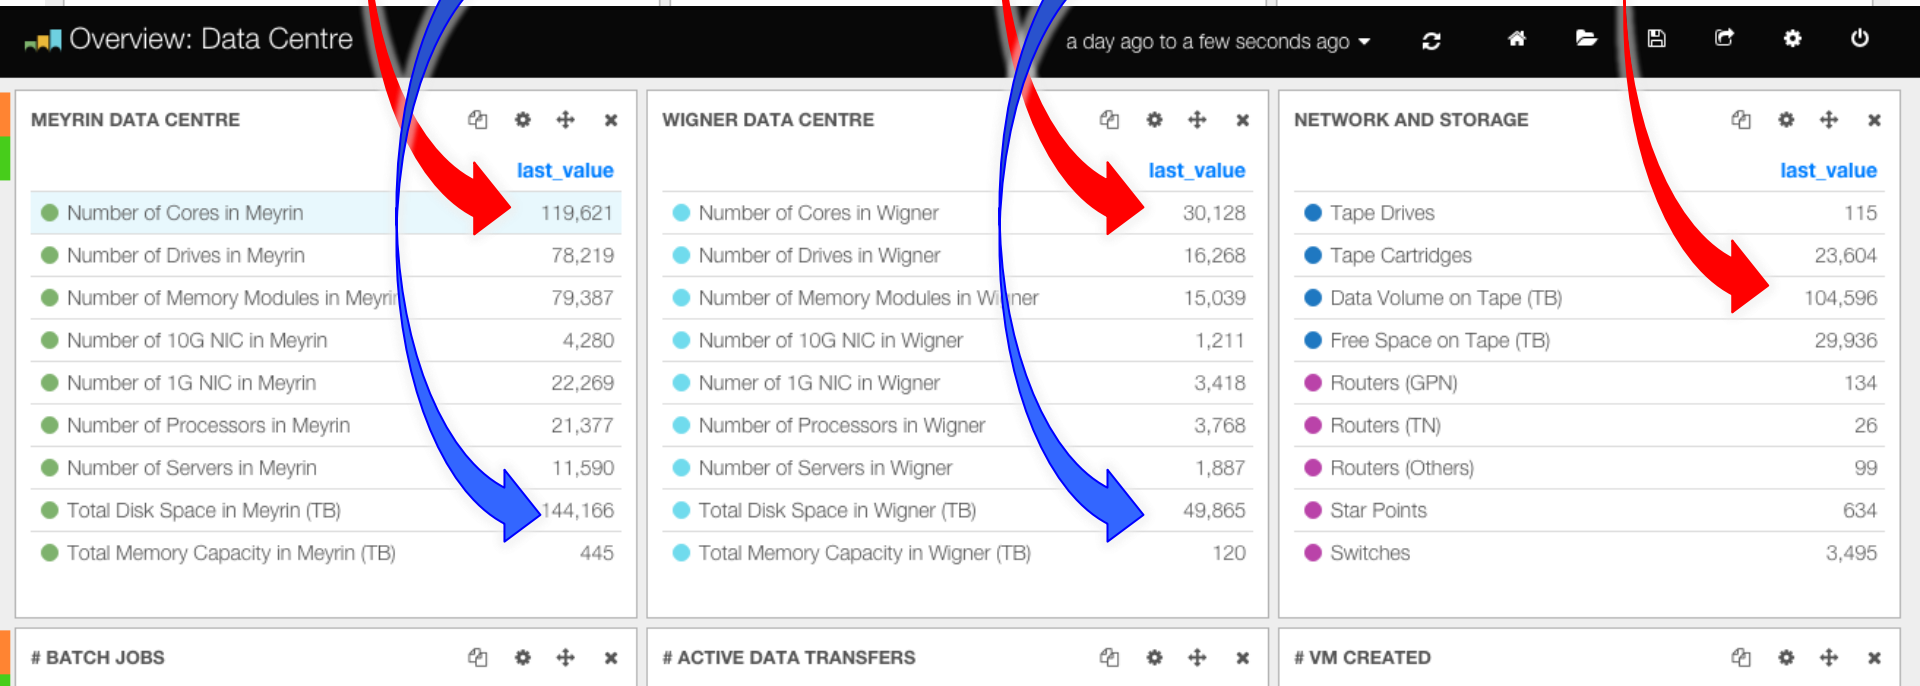

Resource installation

- Mostly on track to be in place for April
- One of two exceptions
  - But expected by ~June
- Not a real problem for 2015, given LHC startup schedule
- Tier 0 had high priority request from AMS for 8000 cores for ~1 month (13 Feb), able to satisfy with just installed Tier 0 capacity

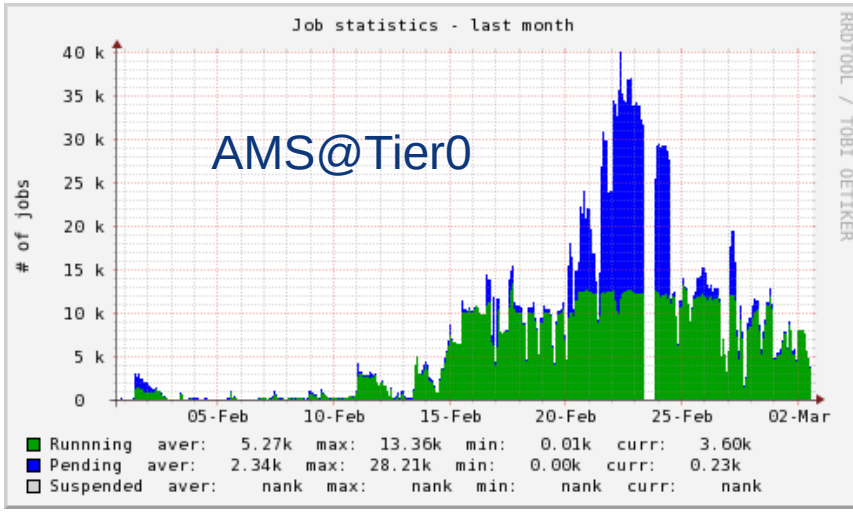

HEP Software Foundation

HEP Software challenge

- Must make more efficient use of modern cores, accelerators, etc.
  - And better use of the memory
- Implies:
  - Multi-threading, parallelism at all levels, optimisation of libraries, redesign of data structures, etc.
- Requires:
  - *significant re-engineering* of frameworks, data structures, algorithms, …
  - *investment of effort* to develop expertise in concurrent programming

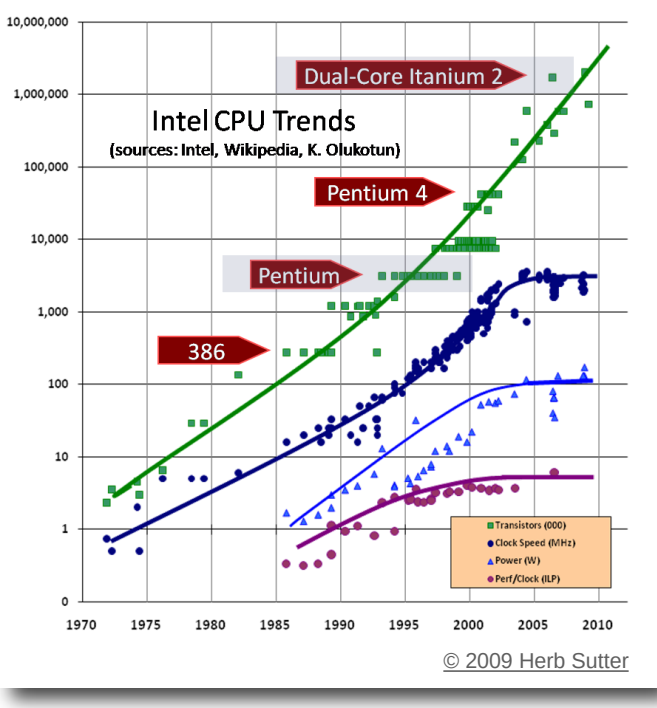

HEP SW foundation: Goals

- Goals of the initiative are to:
  - better meet the rapidly growing needs for simulation, reconstruction and analysis of current and future HEP experiments,
  - further promote the maintenance and development of common software projects and components for use in current and future HEP experiments,
  - enable the emergence of new projects that aim to adapt to new technologies, improve the performance, provide innovative capabilities or reduce the maintenance effort
  - enable potential new collaborators to become involved
  - identify priorities and roadmaps
  - promote collaboration with other scientific and software domains.

HSF - timeline

- April 2014:
  - Initial workshop – to initiate discussions
    - Resulted in a 12 whitepapers to explore what/how
- Summer 2014:
  - Startup team put in place: led by Pere Mato (CERN), Torre Wenaus (BNL)
    - Consolidated whitepaper input – draft of goals etc
    - Set up dissemination – web, mail lists, etc,
    - Gathered a lot of input from interested projects, experiments, etc
- January 2015:
  - 2nd workshop at SLAC
    - Ratify the goals and next steps
- CHEP 2015 (April) Okinawa
  - Opportunity to inform the broader HEP community and encourage participation

HSF SLAC workshop

- January 20-21 at SLAC
  - ~100 people attended (20 remotely)
  - Good representation of HEP labs, experiments, software projects, and other interested communities
- Goal to validate the ideas consolidated by the startup team from the whitepapers
  - And agree a path forward
- Input from other software initiatives
  - Apache software foundation
  - UK Software Sustainability Inst. (esp on training aspects)
  - NSF on building scientific software communities – lessons learned
- Statements from experiments, software projects, institutes
  - All positive and encouraging
  - Many buying into the concept and offering help, or to become guinea pig software projects
  - Common themes:
    - Software knowledge base/catalogue to increase reuse, consultancy for new projects, SWAT teams, build/test infrastructures, teaching, licensing
    - Technical fora, technical discussions with other projects
    - Common software and expertise – avoid reinventing what already exists (e.g. in HPC), help with convergence of solutions and sustainability – community building

Next steps

- Technical forum
  - Place for technology discussion and dissemination of experiences
  - Publish technical notes
  - Help build expertise in the community
  - Concurrency forum – continue as prototype
- Training
  - Consensus that is important initially
    - Several suggestions and volunteers to work with existing schools etc
    - Learn from experience of the UK SSI
    - Working group set up
- Set up SW Knowledge Base
  - Prototype exists
  - Initially try and gather/catalogue software in use and available – provide ability to comment and cross-ref usage
  - Important that community contributes to this
- Build/test/integration infrastructure
  - Mentioned by several groups; examples exist in labs and projects
  - Under consideration
- Under consideration:
  - Licensing issues – must be open source - recommendations needed?
  - Consultancy/SWAT teams – ready to start some activities here – to be better defined and scoped

H2020 project submissions

- EINFRA-7-2014
  - AARC ✔
    - Authentication & Authorization for Research & Collaboration – framework for federated identity platform (eduGAIN)
- EINFRA-1-2014
  - DPINFRA ✗
    - Data preservation services infrastructure, for big-data science
  - EGI-Engage ✔
    - Evolution of EGI
  - INDIGO-DataClouds ✔
    - Building a data/computing platform and tools for science, provisioned over hybrid (public+private) e-infrastructures & clouds
  - RAPIDS ✗
    - Shareable science-domain workflows and services (SaaS) over e-infrastructures to hide complexity; involvement of several EIRO labs
  - ZEPHYR ✗
    - Prototyping & modelling of Zettabyte-Exascale storage systems for future science data

Middleware support

- Concern over lack of support for such key pieces of software
  - ARGUS (security infrastructure)
- Workshop held at the end of 2014
  - Plan for support and development agreed
  - Community support: Nikhef, INFN, EGI
    - Potential new partners
  - This is the model that has been used for DPM etc.

Data privacy

- Changes in EU laws on protection of personal information
  - "user consent" is no longer sufficient
- Need to review and update our AUP and data protection policies
- BUT:
  - Currently some of our (WLCG and experiment) information publication is potentially (now) illegal in many European countries
  - Recent issue affecting xrootd (and sending of monitoring information to sites in US)
    - Technical steps were taken to remove personal information, relocate end-point to CERN (or EU)
  - Potential risk of breaking services if sites are required to stop services, data publication, etc.
- We need to be proactive and address this:
  - New data protection policy being drafted
  - Service developers – middleware *and* experiment – must be aware of this issue and address it correctly
    - Will need to rapidly follow up when problems seen
  - But, sites should also not be too precipitous in reaction –
    - without some discussion of the issue or concern