

HLRmon for advanced visualization of resource usage

Enrico Fattibene (speaker), Tiziana Ferrari, Peter Solagna
INFN

enrico.fattibene<at>cnaf.infn.it
hlrmon<at>lists.infn.it

ISCG 2010 – Taipei March 11th, 2010

Outline

- Motivation
- HLRmon portal
 - Overview
 - Underlying accounting infrastructure
- HLRmon features for WLCG
 - Types of accounting data
 - GUI
- Conclusions and future work

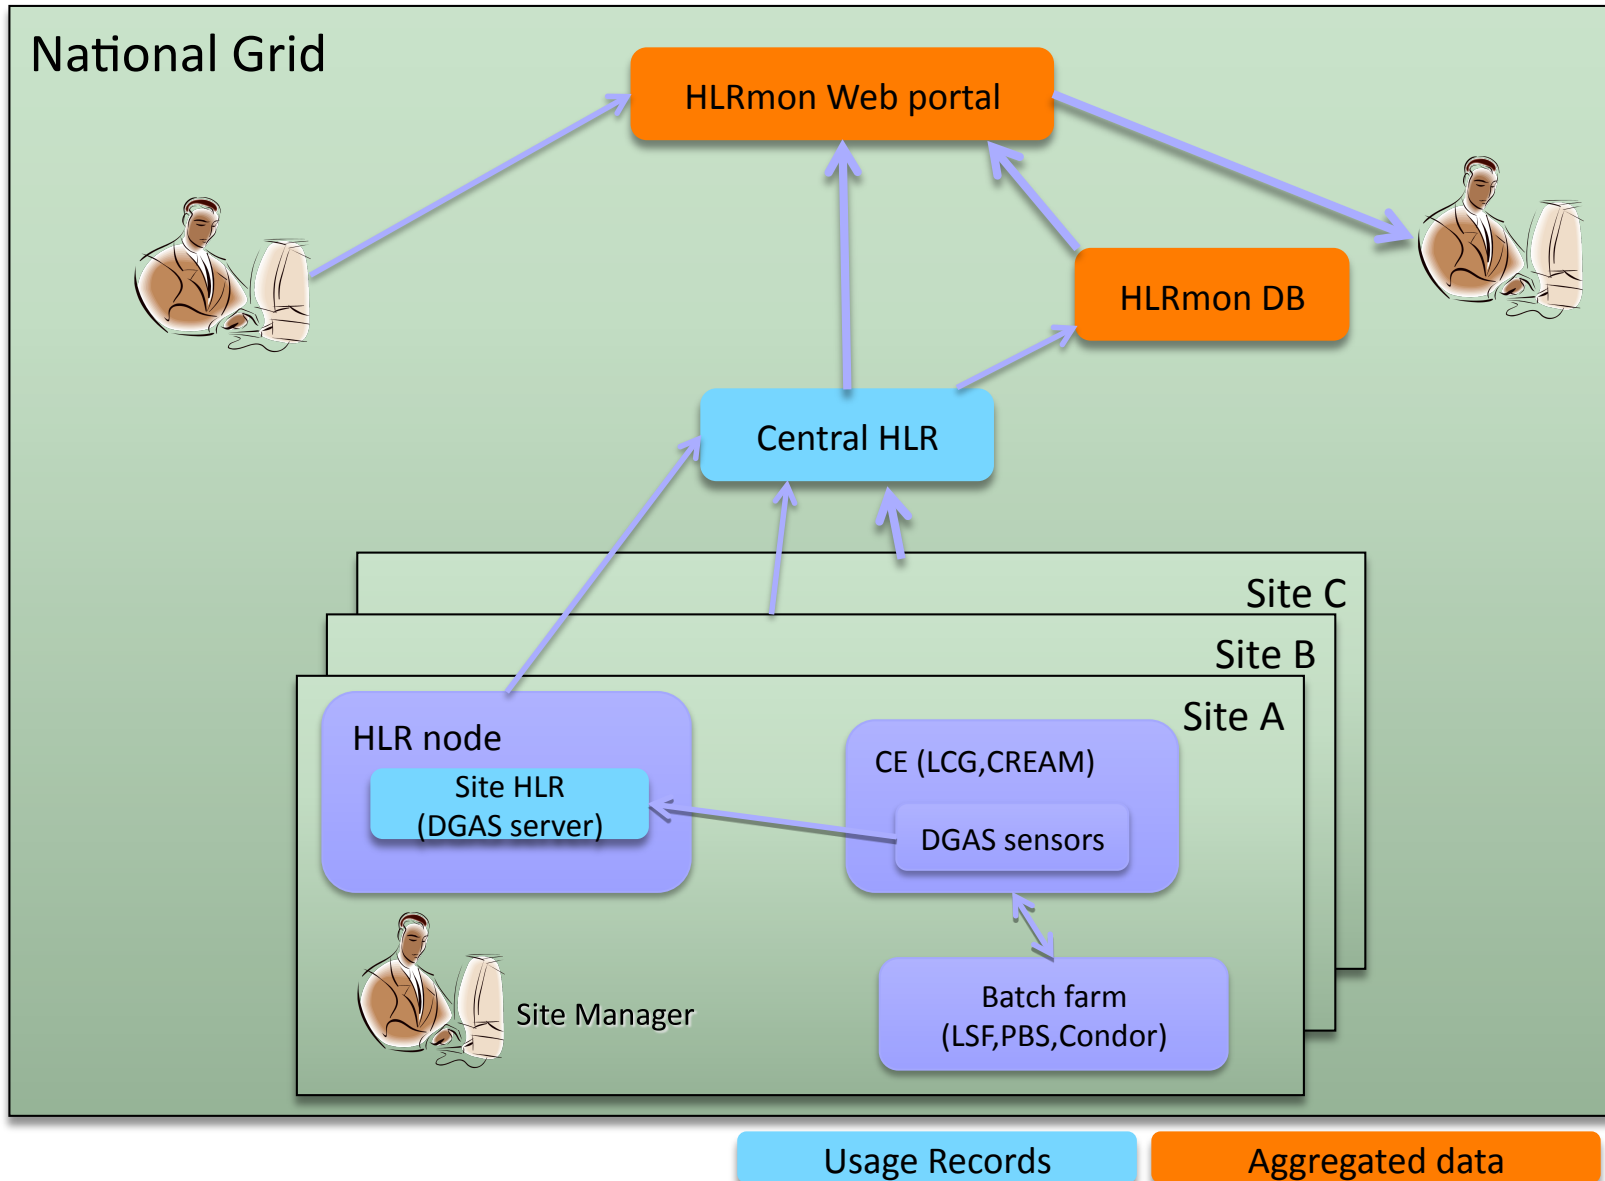
Motivation

- Resource usage accounting in Grid is
 - challenging due to the distributed nature of the infrastructure and the heterogeneous nature of the accessed resources
 - necessary for its stakeholders, such as site managers, VO managers, operation centers, etc.
- Users need to access accounting information at different levels of aggregation
 - an exhaustive and easy to use graphical interface is fundamental
 - many views are needed

HLRmon

- HLRmon (Home Location Register monitor) is the accounting portal of choice of the Italian National Grid Initiative (NGI):
 - Developed and funded in the framework of the INFN Grid project
 - Based on the Distributed Grid Accounting Service (DGAS)
 - Deployed in production since 2007
 - Has been customized to meet the WLCG needs:
 - Specific aggregation views for the Tier-1 and Tier-2 sites and the LHC VOs
 - Both CPU and storage accounting
- <https://dgas.cnaf.infn.it/> (requires a valid certificate released by a CA recognized by LCG)

Deployment scenario



HLRmon: features

- Includes:
 - an open section, which provides aggregated information accessible to the wide public
 - a restricted section providing accounting information per-user: restricted access for privacy reasons
- Different sets of charts
 - CPU and Wall Time expressed in time (days)
 - CPU and Wall Time normalized on the mean power capacity of the site, expressed in kSPECint2000 and HepSPEC-06 and (in the future) other benchmarks
 - Number of executed jobs
 - Job efficiency: CPU Time / Wall Time
- Different keys of data aggregation
 - per site
 - per Virtual Organization
 - per VOMS role
 - per Certification Authority
 - per Job type (Grid vs local)

Why HLRmon?

- Ready to be deployed on a National Grid for the implementation of a stand-alone accounting infrastructure
- Highly customizable views
- Visualization of data at different levels of aggregation
 - site, VO, CA, VOMS role, job type (Grid vs local)
- Visualization of CPU and Wall Time normalized by different benchmarks
- WLCG views which integrate data retrieved by the gLite Information System and storage accounting data
 - Customizable for a desired set of sites and VOs

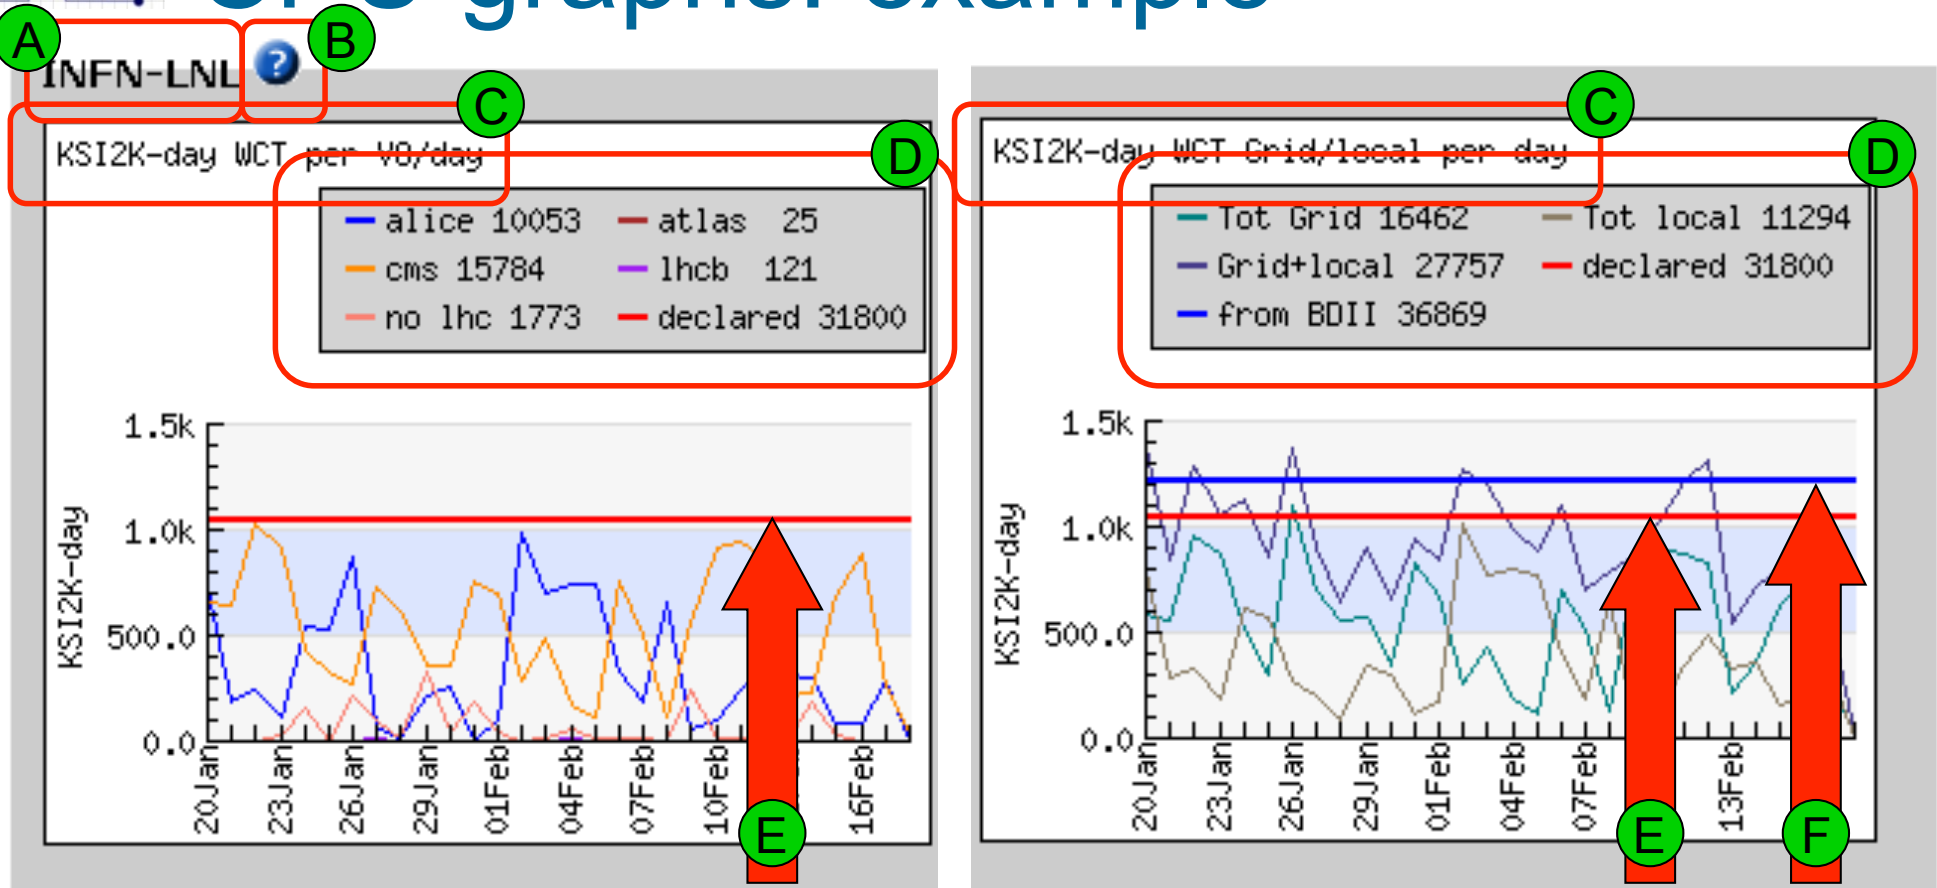
HLRmon views for WLCG

- Why a dedicated set of views for WLCG?
 - To aggregate accounting information per Tier1 and Tier2 sites and LHC VOs
 - To compare resource usage against the WLCG resource pledges
 - To analyse job efficiency at different sites for a given VO, and among different VOs
 - To provide for each site easy access to local fabric monitoring tools (farm, storage, network)

CPU accounting views

- Metrics:
 - CPU and Wall Time, expressed in days and normalized by the mean power capacity for different benchmarks (kSPECint2000 and HepSPEC-06)
 - Number of executed jobs
 - Job efficiency: CPU Time / Wall Time
- Usage levels are graphically compared to the amount of WLCG resources installed for each site, as declared by Grid managers, and to the overall capacity of the site retrieved from the gLite Information System
 - *GlueHostProcessorOtherDescription (Benchmark=HEP-SPEC06)*
 - *GlueHostBenchmarkSI00*
 - *GlueSubClusterLogicalCPUs*
- Types of data aggregation
 - usage levels for LHC and non-LHC VOs
 - Grid vs local jobs
 - Statistics over the last month and the last year

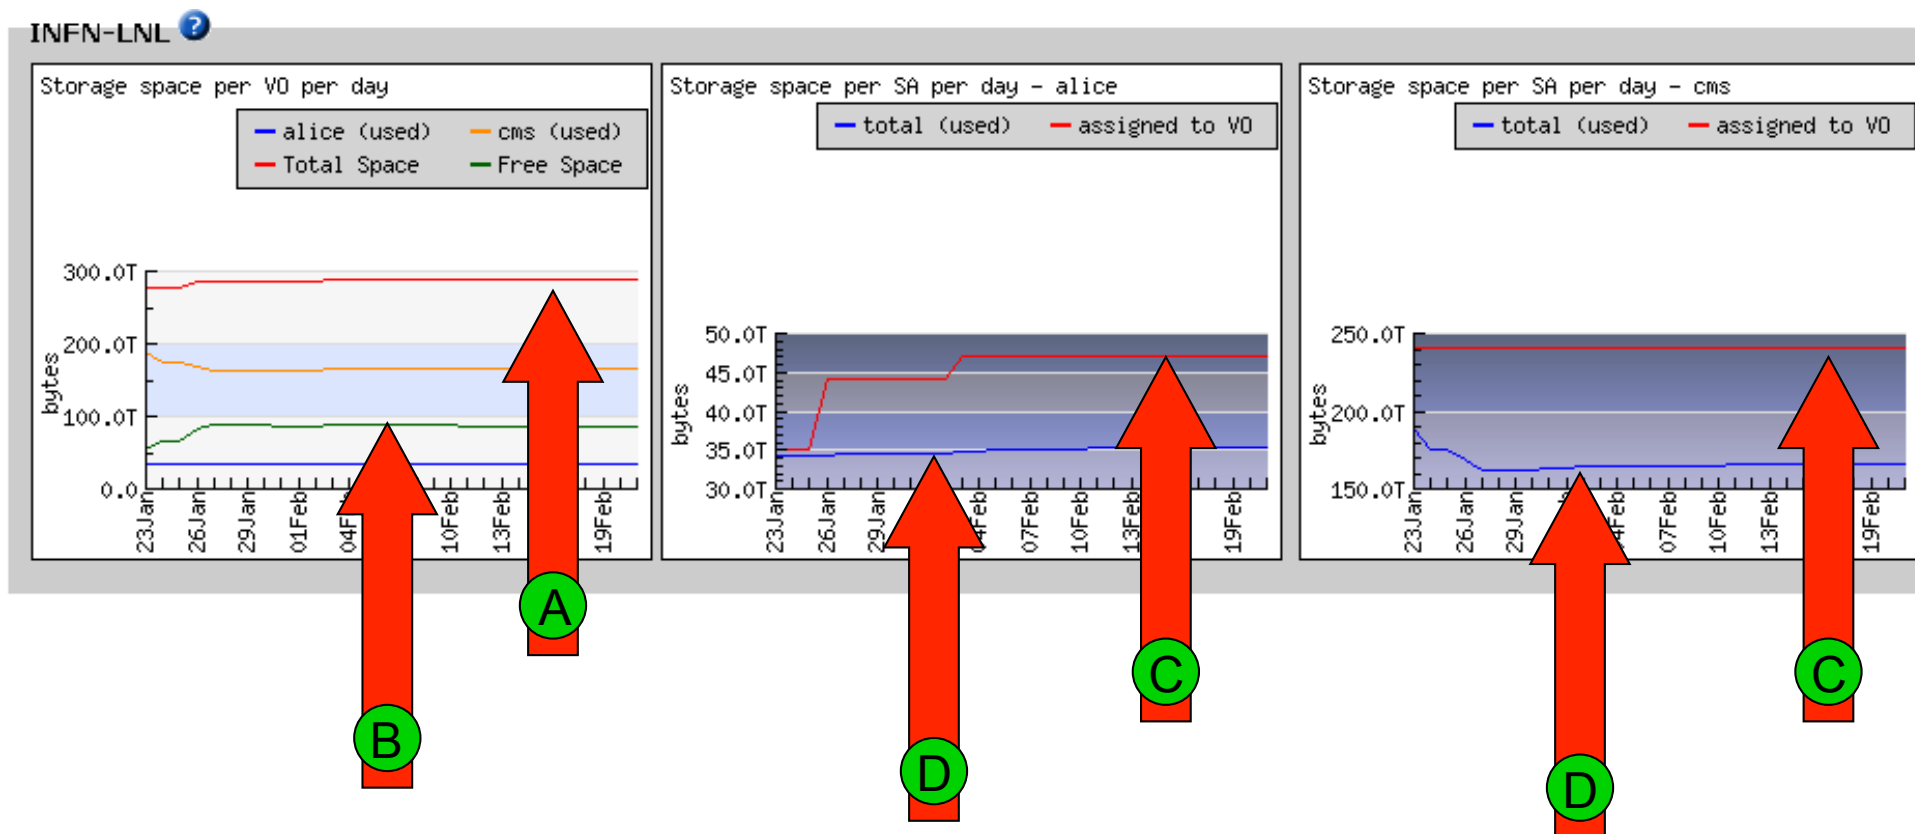
CPU graphs: example



- A - Site name identifying the group of graphs for a given site
- B - Links to local fabric monitoring tools
- C - Chart title detailing the plotted metrics
- D - Legend showing the total metric value computed over the period
- E - WLCG resources as declared by Grid managers (pledges)
- F - Overall capacity (from the gLite Information System)

Storage accounting views

- Metrics:
 - storage space (used, free, total)
 - per LHC experiment
 - per site
 - per Storage Area, where available – currently ATLAS and LHCb
- Probes:
 - ad-hoc probes running daily at each site
 - various SRM implementations supported: CASTOR, dCache, DPM, StoRM
- Monthly and yearly statistics available




A - Total space at the INFN-LNL Tier2

C - ALICE and CMS total space

B - Total free space available in the site

D - ALICE and CMS used space

Storage space per SA per day - atlas



ATLAS storage areas at INFN-T1

Conclusions and future work

- HLRmon is an accounting portal ready for deployment for a national Grid
 - Is based on the DGAS accounting system
 - Offers highly customizable views and different levels of aggregation
 - Includes specific views for WLCG, customizable for an arbitrary set of sites and VOs, providing:
 - storage accounting information
 - easy comparison of accounting data against the WLCG pledged resources
 - RPMS available at:
 - <http://dgas.cnaf.infn.it/hlrmon/hlrmon-web-1.2-1.i386.rpm>
 - <http://dgas.cnaf.infn.it/hlrmon/hlr2sql-1.20-1.i386.rpm>
 - for any information please contact hlrmon@lists.cnaf.infn.it
- Future work
 - Interface to the GOC-DB to retrieve information about downtimes and site status (active or not)
 - Messaging transport of information
 - New views: per-discipline aggregation