

# The “**U**nderstanding **P**erformance!” team in CERN IT

Andrea Valassi (IT-DI-LCG)  
*on behalf of the **iUP!** team*

WLCG Workshop – 9<sup>th</sup> October 2016

# Motivation

- HL-LHC will create a computational problem
  - More data, higher complexity (luminosity, pileup, trigger, detector upgrades...)
  - Gap between computing needs and resources at flat budget
  - ALICE and LHCb will already face a challenge for Run3
- **Long-term activity** needed on ***common performance studies and tools***
  - Link the many parallel (sometimes isolated) activities in the community
  - Aggregate knowledge about existing/new tools for software optimization
  - Shift focus from aggregating to analyzing WLCG monitoring data
  - Aim for quantitative assessment of the performance impact of changes
  - Proof-of-concept implementation of new software and infrastructure
  - High level models of workloads and infrastructures
  - Promote and contribute to common activities between the experiments
  - Play role in WLCG architecture/planning activities
  - *Move cost and efficiency into the spotlight*

# People

- Team in IT
  - Part of WLCG team (IT-DI-LCG)
    - All involved also in other activities, varying % on UP
  - Links with HEP Software Foundation and with the experiments
    - Co-organized performance session of HSF Workshop in May at Orsay
  - Links with Software Technology (formerly Concurrency) Forum and EP-SFT
  - Links with Openlab, Techlab and other groups in CERN IT
- Members
  - Nathalie Rauschmayr
  - Markus Schulz
  - Andrea Sciabà
  - David Smith
  - Andrea Valassi

# Current activities

- Getting an overview of existing activities
  - Within CERN IT and within the community
- Helping to link activities
  - CERN IT Analytics WG, ATLAS workflow performance WG, HSF workshop...
- Developing tools and procedures
  - *By working on concrete experiment applications, workflows and logs*
  - Ranging from macro-level workflows to micro-level software
  - Very brief overview of a few specific examples on the following pages

# Software-related activities

- Analysis of memory usage/dynamics and development of related tools
  - FOM-tools (Find Obsolete Memory), x32-ABI re-evaluation
- Evaluation of tracing tools (SystemTap), studies on CPU hardware counters
  - Aggregate knowledge about other existing tools (Coverity, VTune...)
- Studies on Feedback-Directed Optimization
  - Job profiles helping compilers to auto optimize code (AutoFDO)
- Investigating differences between Intel microarchitectures
  - IPC (instructions per cycle) on Haswell vs Ivy Bridge for some simulation jobs
  - Analyzing effects at the level of instruction caches and branch prediction units
- Analysis of a few specific software components (e.g. Sherpa MC generator)

# Workflow and infrastructure activities

- Analysis of ATLAS jobs on the HLT farms
  - I/O patterns and constraints, workload behavior over time, OS and VM tuning
- Analysis of ATLAS Panda logs (being expanded also to CMS logs)
  - Differences between sites and between job types
  - Focus on understanding workflow efficiency
  - Fit CPU “speed” using data analytics and compare to benchmarks
- Set-up of a small cluster of 16 physical machines for performance studies
  - Batch “standard” nodes, complementary to Openlab/Techlab “exotic” nodes
- Evaluation of computing on storage servers (Andrei Kiryanov)
  - Most WLCG storage servers have low CPU utilization
- Prediction of available network capacity (Hendrik Borrás, Marian Babik)

# Next steps

- Continue with current activities
  - Expand to other experiments (talk to us!) and abstract commonalities
- Identify or build a set of reference “candles” on well defined environments
  - With the goal of allowing comparisons with measurements at grid sites
- Start work on a cost model for a small number of workloads
  - With the goal of predicting quantitative effect of infrastructure changes
- With HSF organize a one-day performance workshop
  - Identify joined projects and build a common roadmap
  - Contribute to the proposed Community White Paper