

WLCG-HSF Systems Performance and Cost Modeling Working Group

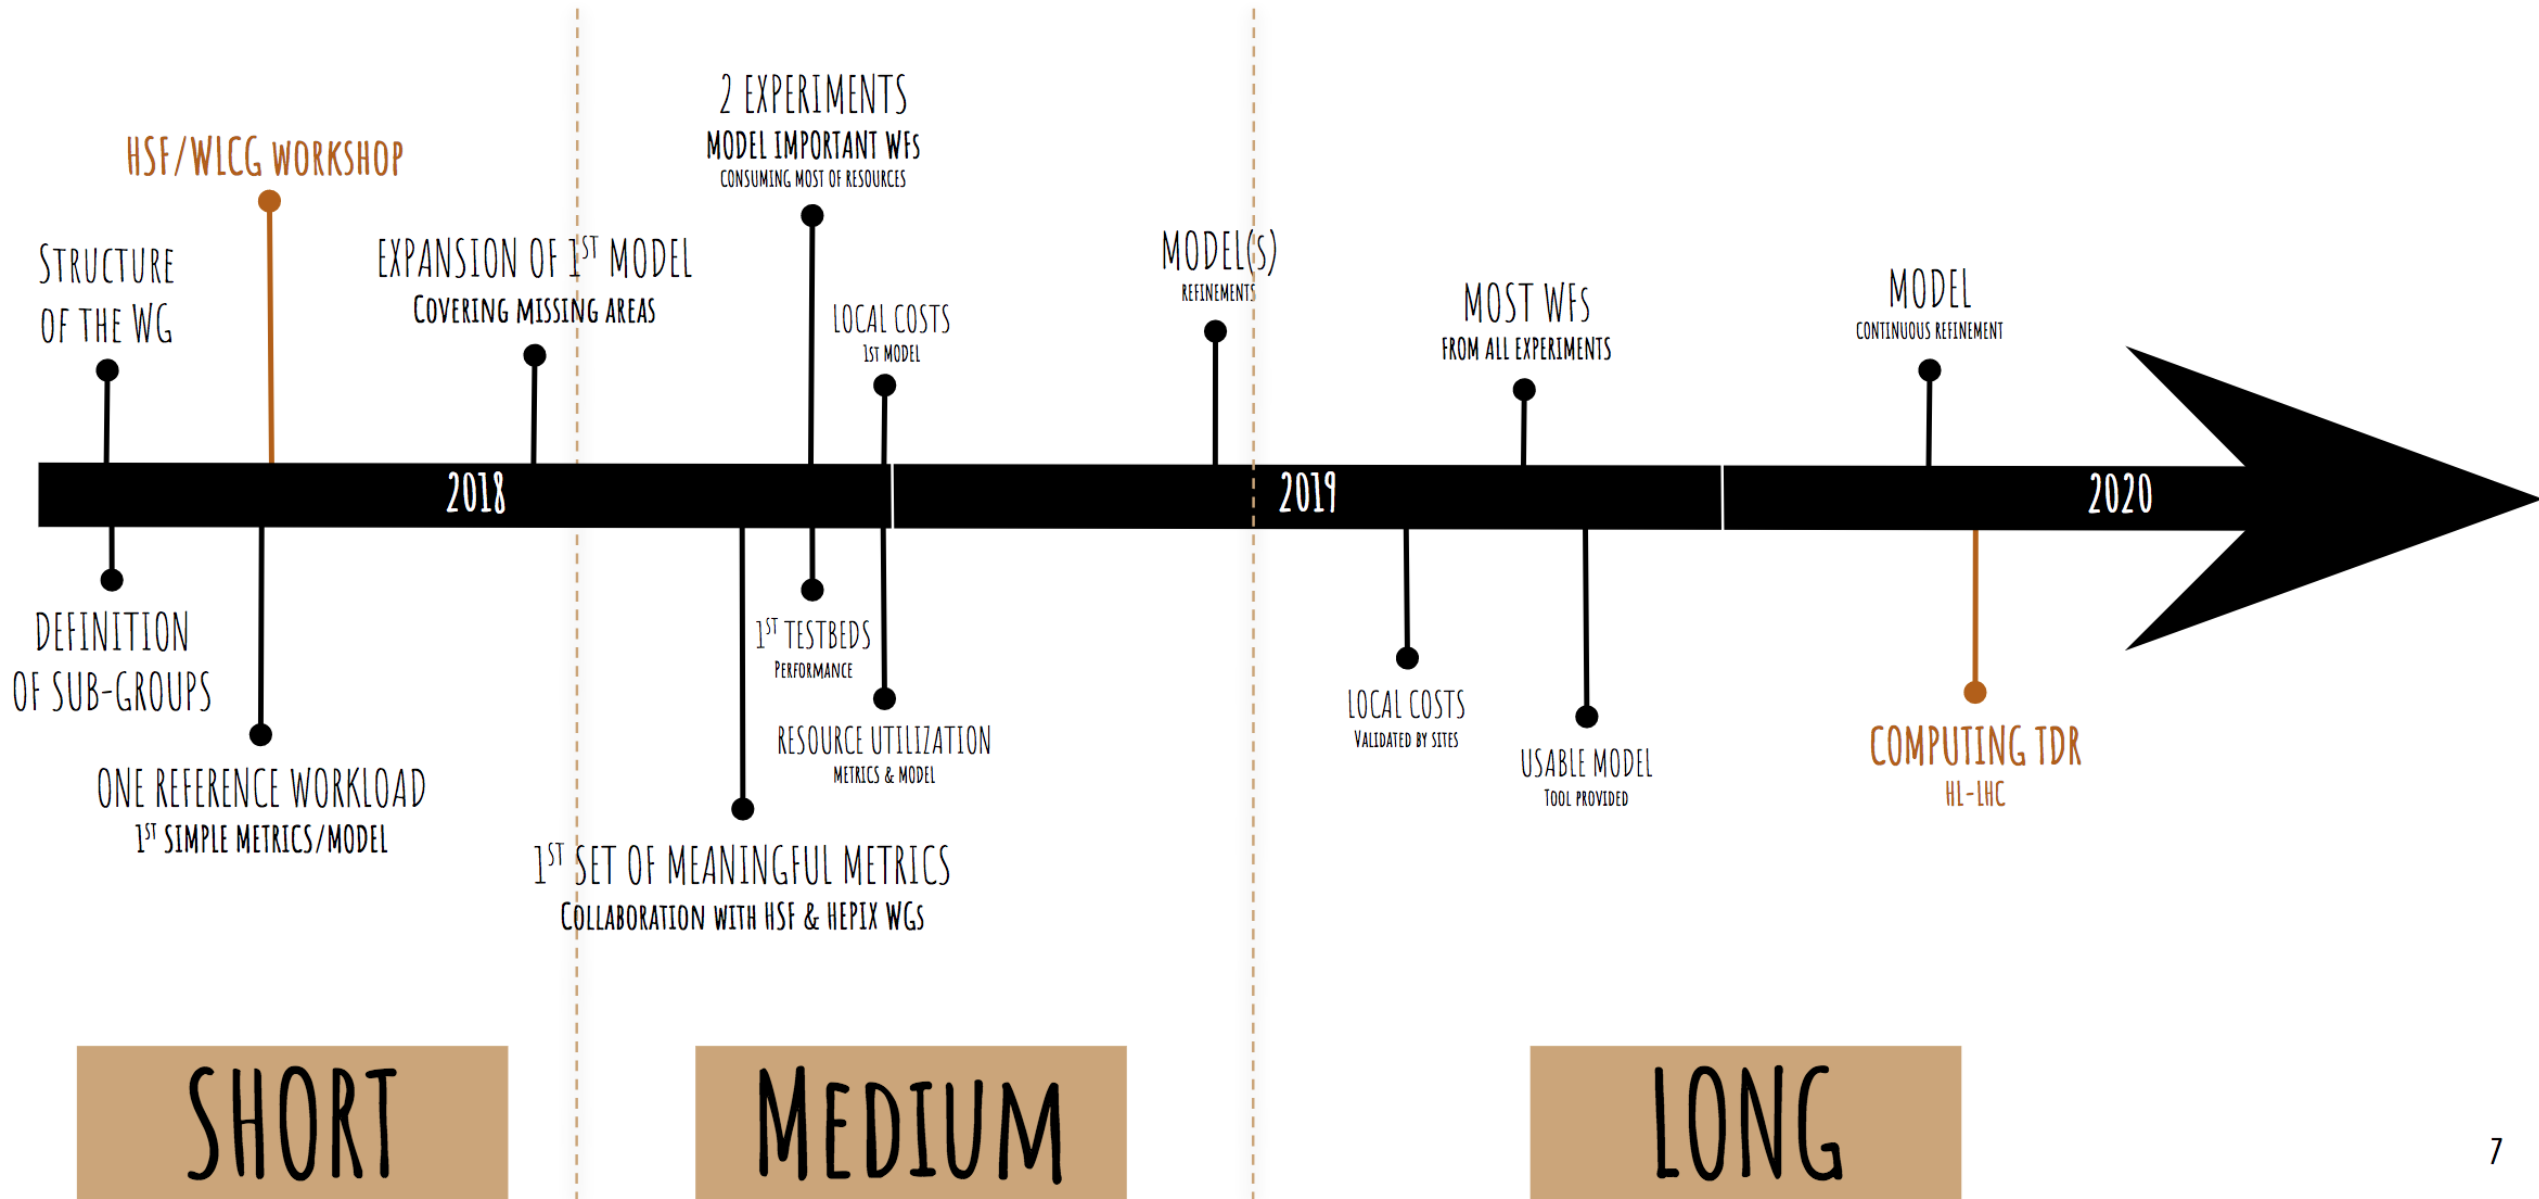
Summary

Markus Schulz (CERN), Jose Flix Molina (PIC), Andrea Sciaba (CERN)

We are a new working group !!!!

- WLCB MB endorsed the creation of the group by 14 November 2017
 - **Joint WLCG HSF working group**
 - 35 active members → wlcg-SystemPerformanceModeling@cern.ch
 - Workload, workflow, framework developers, people who plan, engineer and operate IT sites
- Focus on computing for (HL-)LHC
 - **But starting to understand performance and cost of the current environment**
 - The scope is not limited to WLCG and is of potential interest for other experiments relying on a largely distributed computing infrastructure, like (but not limited to) Belle II or SKA
 - Reporting links between WLCG, HSF and HEPiX

Roadmap [preliminary]



Main Activities:

- Revision of most important workloads for each experiment
 - To create a set of reference workloads accessible to study
 - Common taxonomy
- Packaged versions of the most important workloads
 - To run them easily
- **Definition of properties that best characterise a workload (Metrics)**
 - And measure these
- **Draft a cost evaluation process**
 - Mapping Metrics of workloads to fabrics to local costs
 - Creating a feedback loop to optimize the amount of Physics we can do within budget
- Compile a list of relevant performance analysis tools
 - And make them usable for the average developer/site admin
- Set up a distributed testbed to run tests
 - For laboratory condition measurements
- **Start a simple generic resource calculation model**
 - From spreadsheet to code, transparent

Status

- **All activities have shown considerable progress**
- Some are now straightforward (but still require lots of work)
 - Revision of most important workloads for each experiment
 - Packaged versions of the most important workloads
- Defining metrics moved from definition to building a potential framework for measuring them
 - prmon
 - Stimulating discussions on how gaps could be filled
 - Especially communication from within the workload to external tools
 - Role of batch systems

Status

- Mapping of resource profiles to local costs
 - Exercise with 4 sites showed that there are massive differences what sites consider as cost.
 - Some confusions on goal of a cost model for sites
 - What is it for?
 - How to take into account things that are free...
 - Understanding how many events/year a infrastructure can process given the characteristics of workloads and sites. Guiding purchases and developments.
 - → feedback loop between computing models/workloads/fabrics
 - Importance of complete costs
 - Power infrastructure, racks etc.....
 - Impact of increased complexity on human effort
- A first version of a spreadsheet for calculating costs exists
 - Renaud Vernet (IN2P3)
 - Already useful for sites

Generic resource calculation model

- **Good starting point**
 - Working from Ken Bloom's Python framework
- More parameters from the application metric needs to be added
 - Tuning knobs needed
- **Networks need to added**
- Should evolve into a tool that can answer "What-If" questions to explore the future
- **Current focus: understand what we do today**

Identified (big) Gaps

- Currently we deal with existing workloads and infrastructures
 - Things that exists can be measured relatively easily
- For HL-LHC neither the workloads nor the infrastructures are existing
 - But several ideas exist how they might look
 - New techniques: Machine learning, accelerators,
 - Scale of the computing needs
 - Models: Data Lake, Analysis Farms
 - How to model these ideas ?
 - Some first ideas where discussed
 - Microkernels representing the new approaches
 - Working with the demonstrators/prototypes ...

Next Steps

- Continuation of the current activities
 - They are all on a good path
- Start an activity to explore how we can model HL-LHC conditions
- Trying to recruit more active participation