Ian Bird, CERN

2nd February 2016

Computing model evolution
Introduction and goals

Introduction

- Brainstorming on longer term – HL-LHC era
- LHCC: agreed that by the final LHCC meeting of 2016 we would propose a concrete timetable for working towards an upgrade TDR for HL-LHC by ~2020.
- May seem that this aims mainly at ATLAS and CMS, since the planning for LHCb and ALICE for the phase 1 upgrades are more advanced.
  - However, we should really aim to be as inclusive as possible and indeed to have an eye on the broader HEP and global science communities as we think about this.

Assumptions and boundaries

- The current guidance for Run 2 and the foreseeable future is that we should plan around flat budgets for computing. There is presently no indication that this situation is likely to change on any timescale, but with some potential concern that it might decrease.
  - However, we should first understand, and state, what we think the cost of HL-LHC computing will be to do the advertised physics
- Anticipated data volumes seem to be on the order of 0.5-1 Exabyte of raw data per year combined across 4 experiments, without major changes in philosophy of data selection/filtering.
- Compute load with the naïve extrapolations so far made, have computing needs between 5-10 times in excess of what may be affordable as anticipated by technology evolution (which of course itself is very uncertain).
- Given these two conditions, we must plan for innovative changes in the computing models from DAQ to final analysis, in order to optimise the physics output.

HL-LHC computing cost parameters

New grid/cloud models; optimize CPU/disk/network; economies of scale via clouds, joint procurements etc.

Business of the experiments: amount of Raw data, thresholds, but also reconstruction, and simulation algorithms
Cost of simulation/processing should be a design criterion of detectors

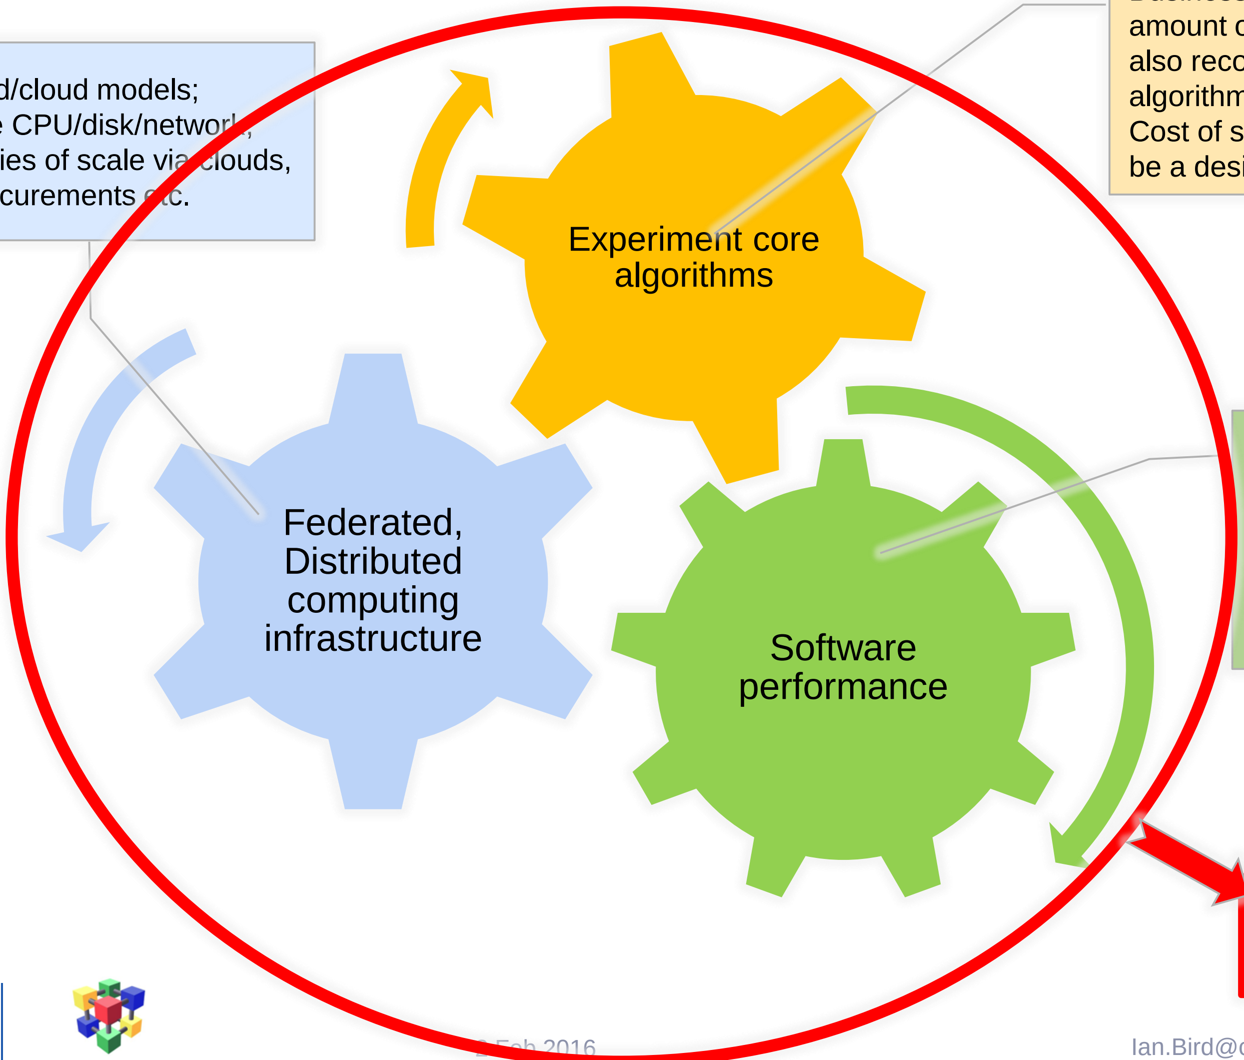

Performance/architectures/memory etc.;
Tools to support: automated build/validation
Collaboration with externals – via HSF

COST OF COMPUTING

Cost drivers – 1

- The amount of raw data to be processed. Can this be more intelligent – e.g. by moving the reconstruction into the “online” (which does not necessarily have to be at the pit – but make a decision based on accurate reconstruction before recording)?
  - Actually is it even realistic now to assume that we can blindly write raw data at a maximum rate and assume it can be sorted out later?
  - Storage cost reduction with compression (lossless and lossy) – this will not reduce the offline load, but may help manage the flow.
  - Clearly most of this implies in-line automated calibrations and alignments, and single pass high quality reconstruction.
  - Ultimately data volume management might require physics choices.
  - Consideration of computing costs (esp for simulation) during the design of detectors for the upgrades. (e.g. ATLAS calorimeter is expensive to simulate).

Cost drivers – 2

- New algorithms and re-thinking of how reconstruction and simulation are performed.
  - This area needs some inspired thinking and innovation around algorithms.
  - Adaptation and optimisation of code for different architectures which will continually evolve.
  - Rethinking of how data & I/O is managed during processing to ensure optimal use of the available processing power.
  - All of this probably requires a continual process of automation of build systems and physics validation, since it is highly likely that key algorithms must be continually optimised across architectures and compilers. The current situation of single executable for all architectures will not be optimal.

Cost drivers – 3

- Optimisation of infrastructure cost. Many parameters that could be investigated here.
  - Reduce the amount of disk we need, to be able to acquire more CPU/tape. Today disk is the largest single cost, and is mostly driven by the number of replicas of data.
  - Can we get economies of scale through joint procurements, large scale cloud providers, etc. On this timescale we may expect commercial computing costs (of CPU and cache) may be much cheaper that our internal costs, if we can reach a large enough scale.
  - Likely we will have significant network bandwidths (1-10 Tb/s) at reasonable commercial costs.
  - Should think about very different distributed computing models – for example a large scale (virtual) data centre integrating our large centres and commercially procured resources. Data would not be moved out of that “centre” but rather processed in-situ and available remotely to the physics community. Such a model would minimise storage costs, and enable potentially different analysis models since all of the data could be (virtually) co-located. Simulation load would be maintained at remote centres.

Goal

- Assume we need to save factor 10 in cost over what we may expect from Moore's law

- 1/3 from reducing infrastructure cost
- 1/3 from software performance (better use of clock cycles, accelerators, etc. etc)
- 1/3 from more intelligence – write less data, move processing closer to experiment (keep less) - writing lots of data is not a goal

Other considerations

- In this timescale, computing (and thus physics) is going to be limited by cost, thus we must make every effort to maximise what can be done,
  - in terms of infrastructure,
  - but also where it costs real skilled effort (such as algorithms and performance optimisation).
  - This requires cross-boundary collaboration (experiments, countries, etc.)
- There is significant scope here for common projects and developments, common libraries, frameworks, etc.
  - A significant level of commonality is going to be required by our reviewers (scientific and funding).
  - Four solutions are not going to be sustainable, but we must justify where differences need to be made.

Radical changes?

- Move processing that produces science data (for physics analysis) e.g. AOD’s or ntuples(?) as close to DAQ as possible before distribution
  - Serious cost benefit, but serious political problem to convince FA’s (this is why we have a grid now)
  - More or less the LHCb and ALICE models for Run 3
- Limit how much data we ship around (to limit disk costs – today 2/3 of budget is disk)
  - The idea of a “virtual data centre” where the bulk data is processed and stored
  - Potentially a way to get economies of scale, etc

- DC primarily academic, allow hybrid
  - Data replicated across all
- Compute: commercial or private – cost/connectivity
- Data input from external sources: instruments, simulation factories
- Clients see as a single DC, access data products: existing, on-demand
  - Allows new types of analysis models

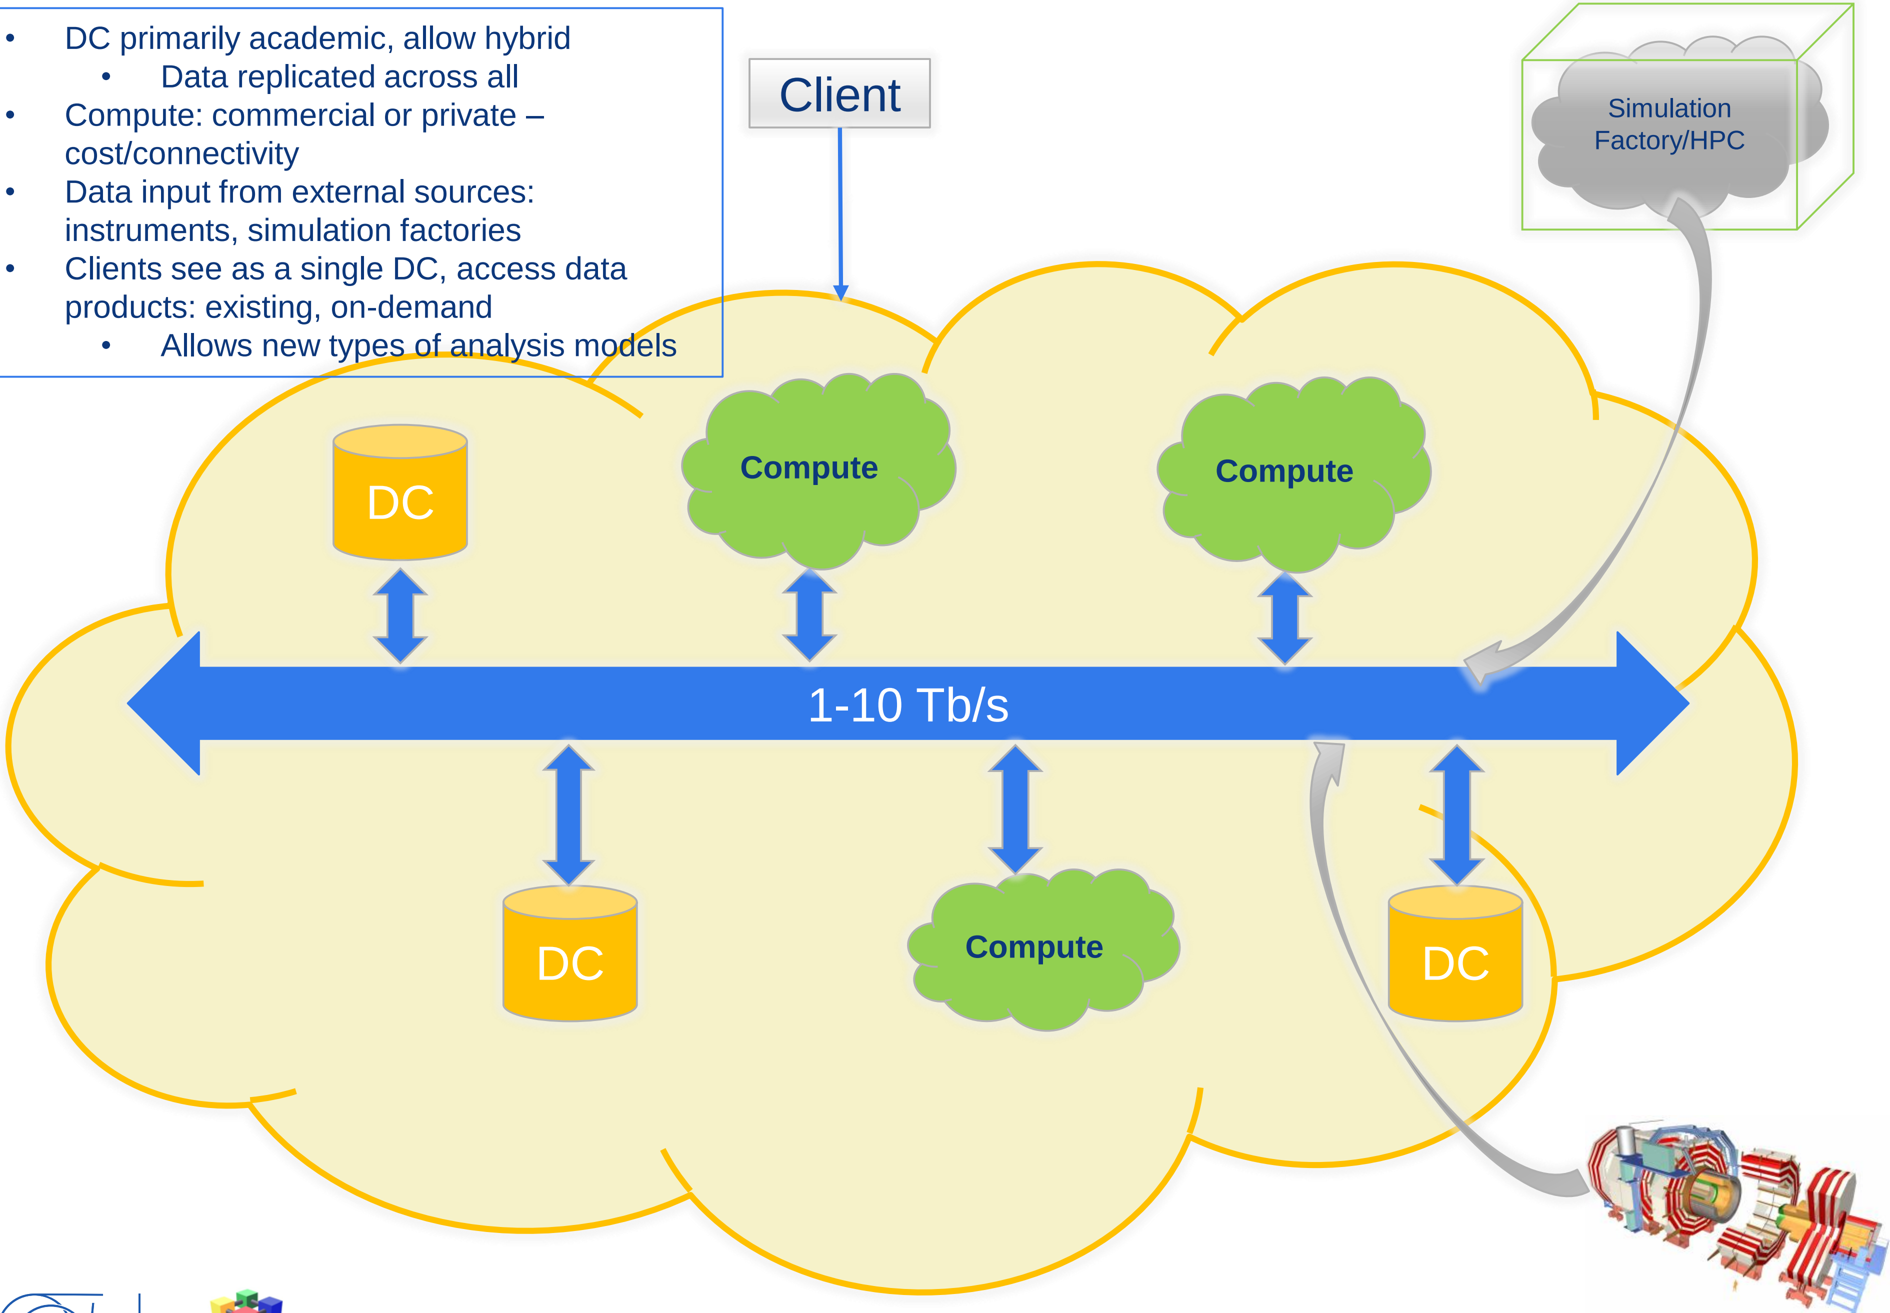

A Virtual Data Centre for Science?

Zephyr

- Was a (failed) proposal to the EC last year to build a distributed data infrastructure along these lines
  - Was to develop a prototype “Zettabyte” data infrastructure (EC call)
  - Essentially a “RAID” between data centres (storage systems), along lines of previous slide
  - Policy driven replication of data between centres
  - Model allows to use commercial DC’s in a ”transparent” way if desired
- Some interest expressed by OB in building such a prototype anyway with Tier 1s

HEP Facility timescale

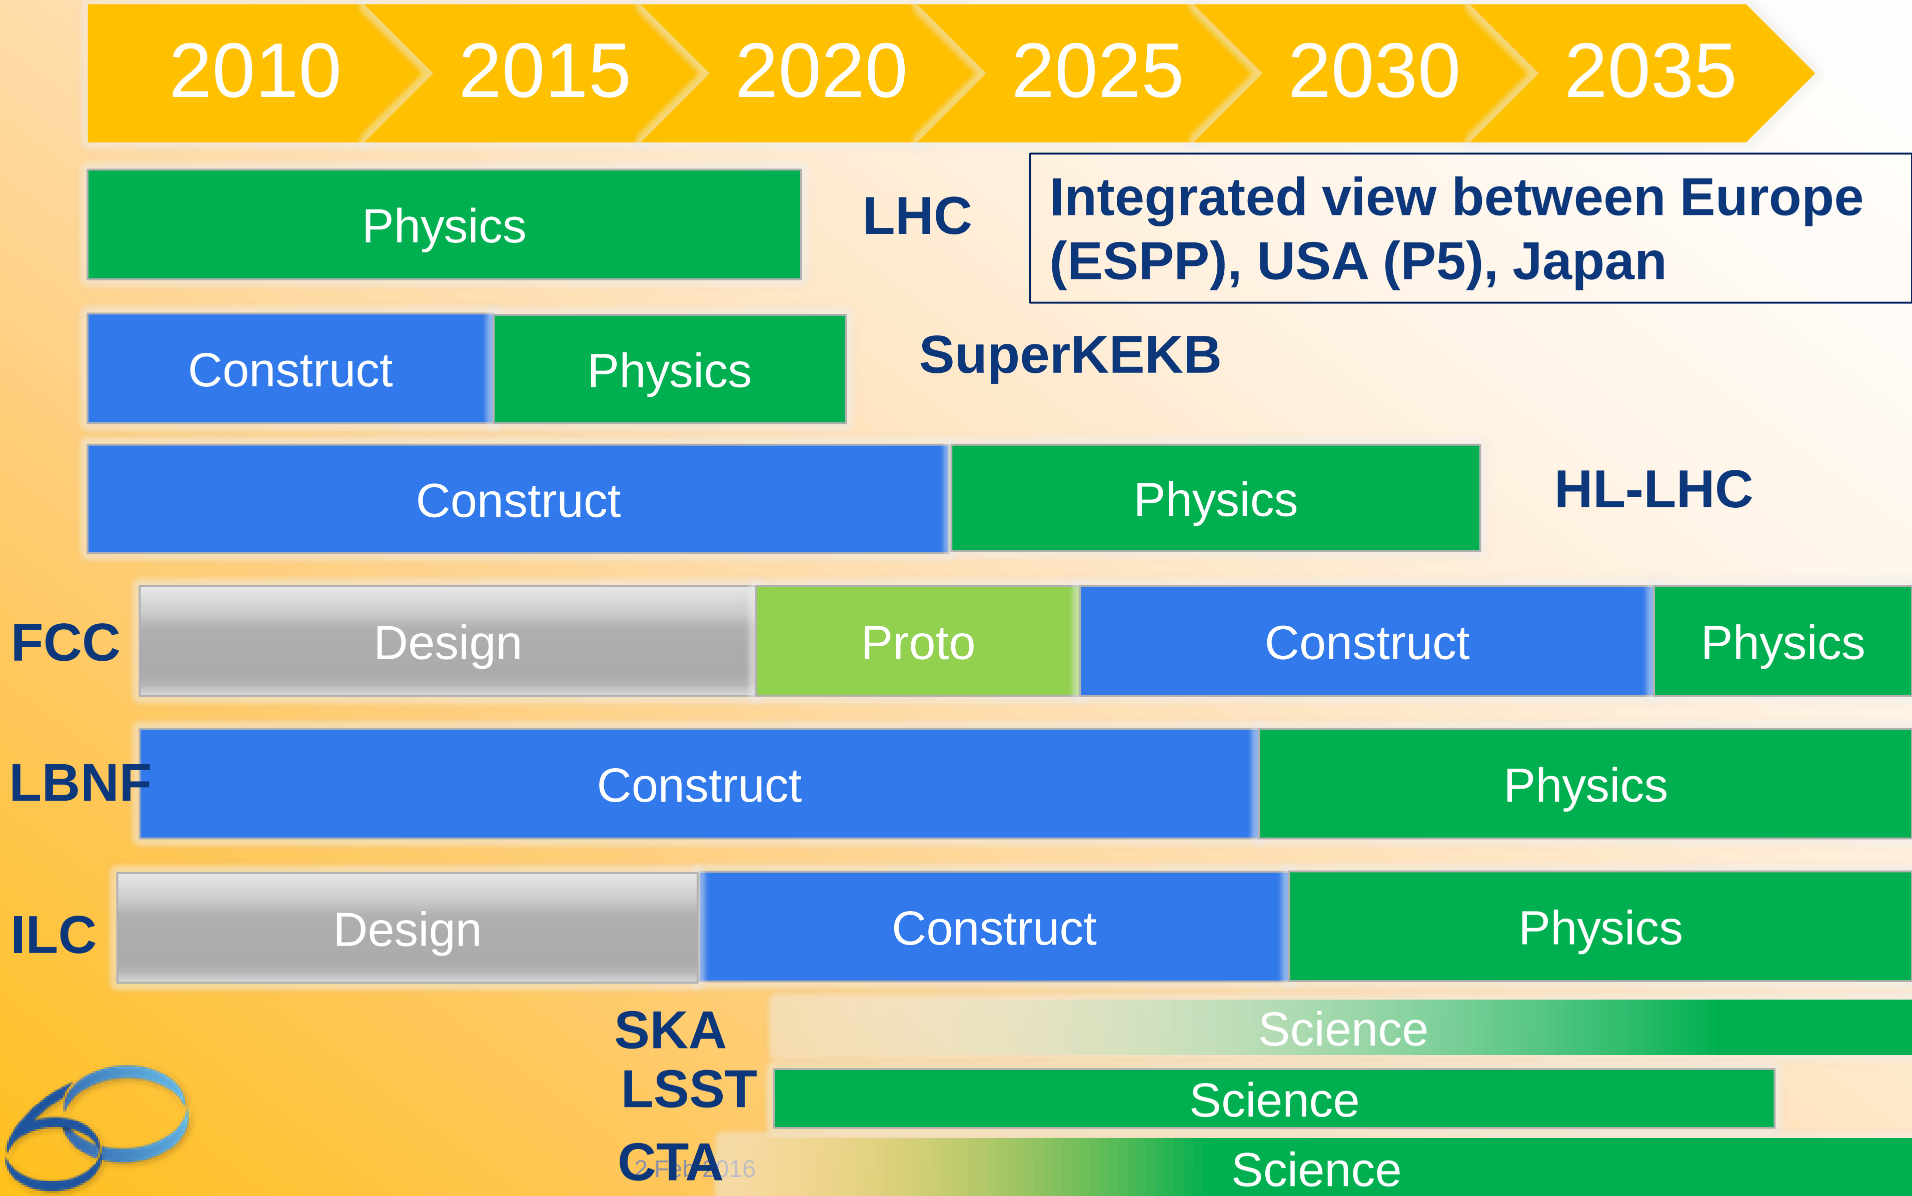

Integrated view between Europe (ESPP), USA (P5), Japan

Longer term

- Should WLCG evolve to a broader scope for HEP in general that covers the evolution and needs of planned and future facilities and experiments?
- Other astro(-particle) experiments will share many facilities of HEP
  - Need to ensure that facilities don't have to support many different infrastructures

Today's goals

- Open brainstorming - seeded by questions and brief inputs
  - Not answers! → ideas and suggestions
  - Feel free to provide input – e.g. 1 slide of ideas
- Will need to be followed up after the meeting
- Would like to have a document that
  - Outlines areas of R&D work that we need to do
  - Perhaps 2-3 strawman models of what computing models could look like in 10 years

Agenda topics

- Initial thoughts from experiments on main areas of concern
- Infrastructure models – what could we do?
- Workflows and tools – what convergence is possible?
- Technology evolution – how to track/predict?
  - How to support HSF effort
- Modelling of infrastructure/sites etc – we need something to test ideas (at gross level), can we build a cost model?
- Software (tomorrow) – commonalities, how can HSF help, etc.?

Timeline

- Need fairly quickly (~1 month) outline of areas of work and R&D effort
  - Will discuss outcome of this workshop with LHCC in March
- By end of 2016 need to have a timeline that leads to a TDR by ~2020
  - Timeline of R&D, testing, prototyping etc
- Hopefully good ideas can be deployed earlier!