Ian Bird
Scientific Computing Forum
CERN
15th February 2017

Scientific Computing Strategy (for HEP)

Preamble & disclaimer

- When LHC Computing started there were no large Internet companies (Google, Amazon, …)
- HEP had to build infrastructure and tools
- Starting today might be very different …

- The following is intended to provoke discussion - it is not a plan

Some drivers from HL-LHC

- Need to fit computing within constrained costs
  - But being asked to do much more …
- Often need to co-locate (intimately) with other sciences in our Data Centres
  - Esp. with Astronomy/Astro-particle & other physics e.g. photon science etc

Some consequences

- Need to constrain costs
  - The main driver for WLCG is overall data volume (replication) & wide distribution
    - 2/3 of total global cost is in disk
- Infrastructure must no longer be (too) special
- Need to be able to use commercial & opportunistic resources
  - Including non-traditional for HTC: HPC, cloud, special architectures etc.
- Thus need significant agility and performance optimisation through software
- Learn from our experience and that of large internet companies
- Need flexibility/agility to changing markets – e.g. cost of commercial resources, obsoleted technologies (perhaps overnight)
- Must recognise and leverage opportunities of commonality at all levels
  - Between experiments, HEP, across disciplines, with industry, …

➢ Failure to change will limit the scientific output

HEP Facility timescale

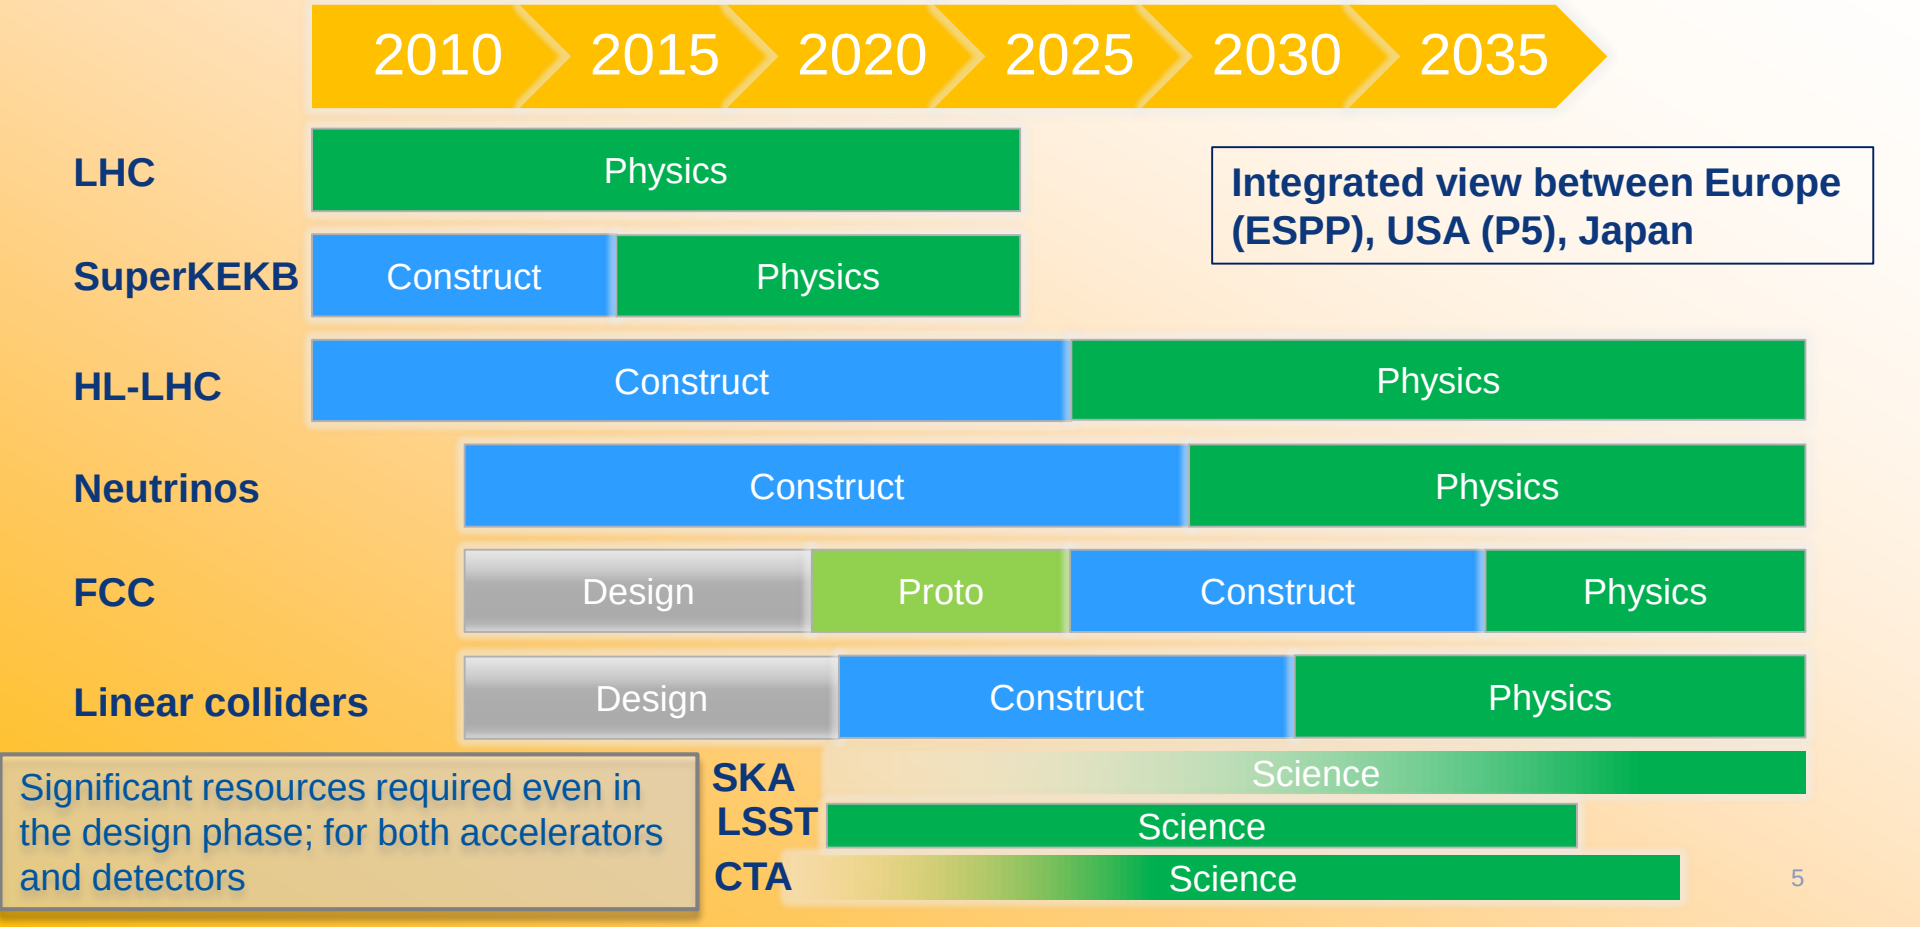

Integrated view between Europe (ESPP), USA (P5), Japan

Significant resources required even in the design phase; for both accelerators and detectors

HEP Computing over 10 years

- HL-LHC is the major challenge
  - But also: neutrinos, Belle-II, development for linear collider experiments, FCC, etc.
    - And HENP facilities will also host Astro/Astroparticle experiments
- WLCG is a global collaboration, and successful in supporting LHC computing
  - Many lessons learned:
    - Technology, sociology, funding
- Can we imagine a HEP-wide computing environment?
  - Building on what we have, evolving and providing tools and infrastructure services to be used or adapted to future needs of HENP

Key Components

- A general infrastructure and services (data stores, compute facilities, networking, etc.);
  - Associated services like AAA, security, base monitoring, operational support
  - Needs to be capable of supporting different computing models and agile to technology changes
- "Software": full stack from workflow and data management tools to application level; Common tools, libraries, etc.;
  - BUT: a set of optional tools, contributed, developed, maintained, by the community;
  - Common R&D and support tools (technology tracking, software tooling);
  - This is essentially what the HSF is mandated to do.
- A Steering group structure to organize and evolve the above – mandated by e.g. ICFA
  - Constituted from major global HENP facilities, experiments, observers (A-P)
- Experiment/facility-specific (optional) processes
  - e.g. WLCG for the LHC resource management, etc.

Infrastructure

Infrastructure

- Networks
  - LHCOne – global HENP (and A-P) overlay network (LHCOPN is LHC-specific),
  - Address the need for access to commercial clouds via NREN's
  - How to manage a HEP data cloud network (multi-Tb interconnect, SDN)
- AAA services
  - Federated identities, authorization mechanisms, accounting
  - Policies
- Operational services
  - Support, incident response, monitoring, etc.

Reducing the cost of data (HL-LHC)

- Decrease the data volumes (large volumes are not the goal)
  - More and early filtering?
  - Reduce data formats?
  - Do not keep RAW as a primary data set?
- Distribute less
  - Distribution vs caching (costs are different)
- Make more use of high-latency/low-cost storage (“tape”)
- Buy less disk

Possible Model for future HL-LHC computing infrastructure

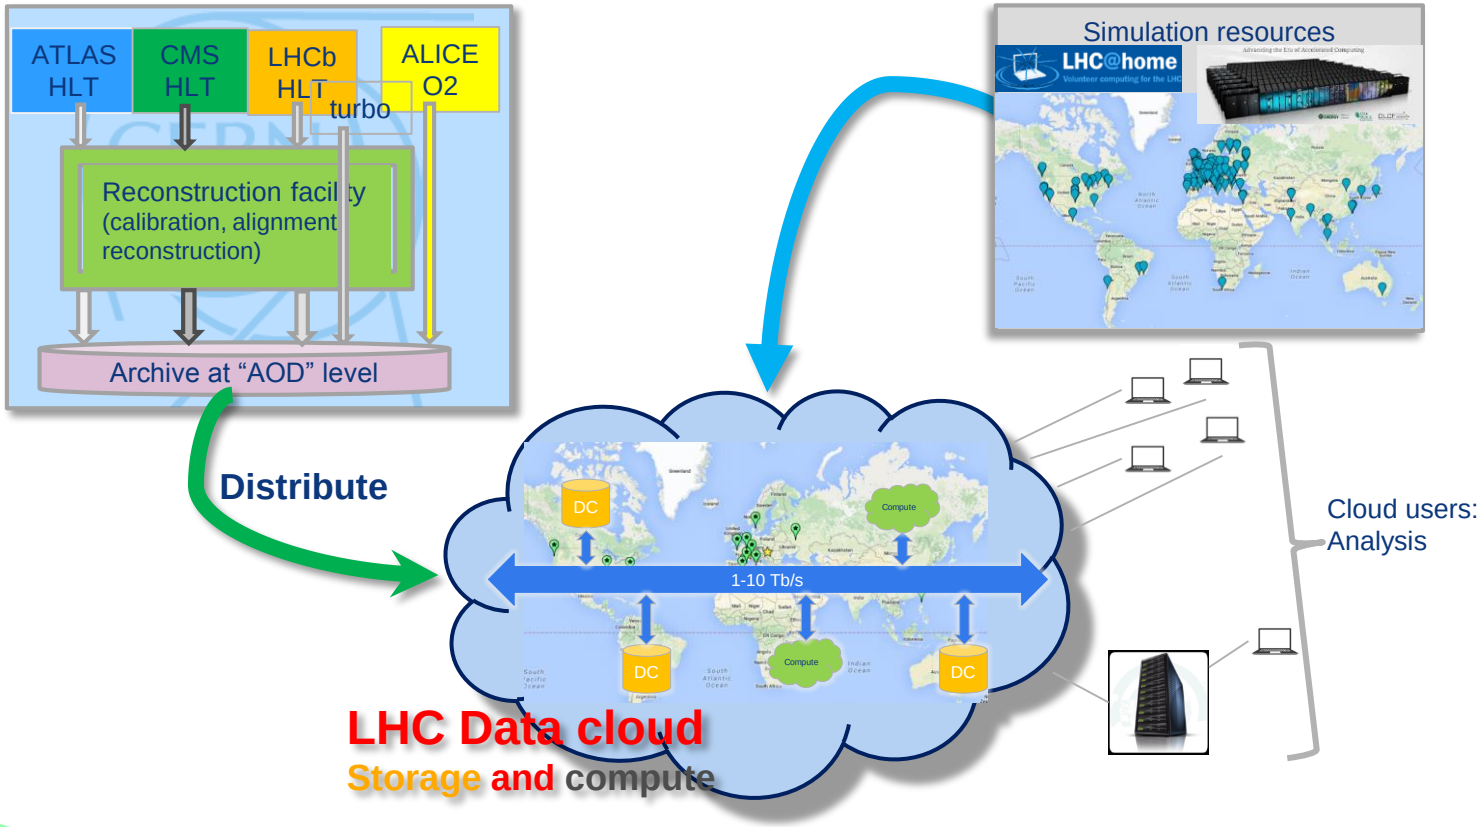

Straw-man for HL-LHC (and …)

- Production of analysis-ready data sets should be done near-line
  - Integrate online & offline
  - Real time calibrations/alignment etc
  - 1-pass (& final) reconstruction
  - Produce “AOD”-level data sets
  - ? Never distribute raw data, only AOD-level
    - ? Keep only 1 copy of raw (or none!)
- ☞ Need sufficient local computing to be able to do this; “Wigner-like” or very remote is not suitable

LHCb turbo, ALICE O2;

CMS & ATLAS start to go in this direction

Build a “data cloud”

- Few – O(5-10) - large centres
  - Multi-Tb private (SDN) network between them
  - Treat as a single “virtual data centre”
    - Policy replicates data inside for security and performance
    - Think of RAID across data centres
  - Store all of the “AOD” data here; Do not replicate data to global physics institutes (major cost)
- Pluggable compute capacity:
  - HEP resources at these centres & other large centres
  - Commercial compute
- Model allows commercial data centres
  - For storage – enough redundancy that a commercial centre could unplug
  - For compute
  - Relies on networking and GEANT/Esnet etc. connections to commercial entities, policy
- ☞ Users access data in this cloud remotely
  - Eventually download “ntuples” – or equivalent
  - All organised processing is done in this model
- ☞ Enables new analysis models: all data can be seen as colocated
  - Get away from the “event-loop” → queries, machine-learning, etc.

This idea has been discussed in the WLCG community (e.g. see I. Fisk CHEP plenary)

- Hybrid model:
  - HEP-resources at a level we guarantee to fill → cost-effective
  - Commercial resources for “elasticity”
- Needs new funding models

What about Tier 2s, opportunistic, … ??

- Depends on scale
- Large Tier 2s & HPC could be pluggable compute facilities;
  - need connectivity to the private data network
- Most effective use of medium and small centres is for simulation
  - Simulation loads are huge (~50%)
  - Centres upload MC samples to the cloud
  - Ideally want full-chain pre-emptable simulation workflows
    - To make best use of HPC, opportunistic, volunteer, etc
- Some HPC potentially best suited to event generation
  - Becomes a significant compute load
- Must consider the best way to use what we have, not focus on the labels/prestige etc.

Easily adapts to the situation where a country may replace academic computing clusters with “tokens” for commercial cloud providers

Other comments

- This virtual data cloud model may be very interesting for other sciences
  - E.g. SKA Regional Centres
  - Works also for DUNE, Future facility development, others
  - Can provide resiliency and long term preservation capabilities
  - Scale-out is inherent;
  - Model is very much like large commercial cloud providers
- Case for new CERN Prévessin DC
  - even without HLT farms; although synergy with them co-located is potentially significant:
  - ➔ should see it as an extension of the DAQ to produce analysis data
- Requires (potentially) significant changes in funding models
  - Can we actually procure commercial resources at large-enough scale to get economy?
    - HNSciCloud as a proof-of-principle of joint procurement
    - Can we purchase from the largest cloud vendors? Politics?
    - Real cost-efficiency and elasticity requires a “spot-market” price

Software

Software

HSF Set up in response to recognition that software will be key to success for HL-LHC and the future

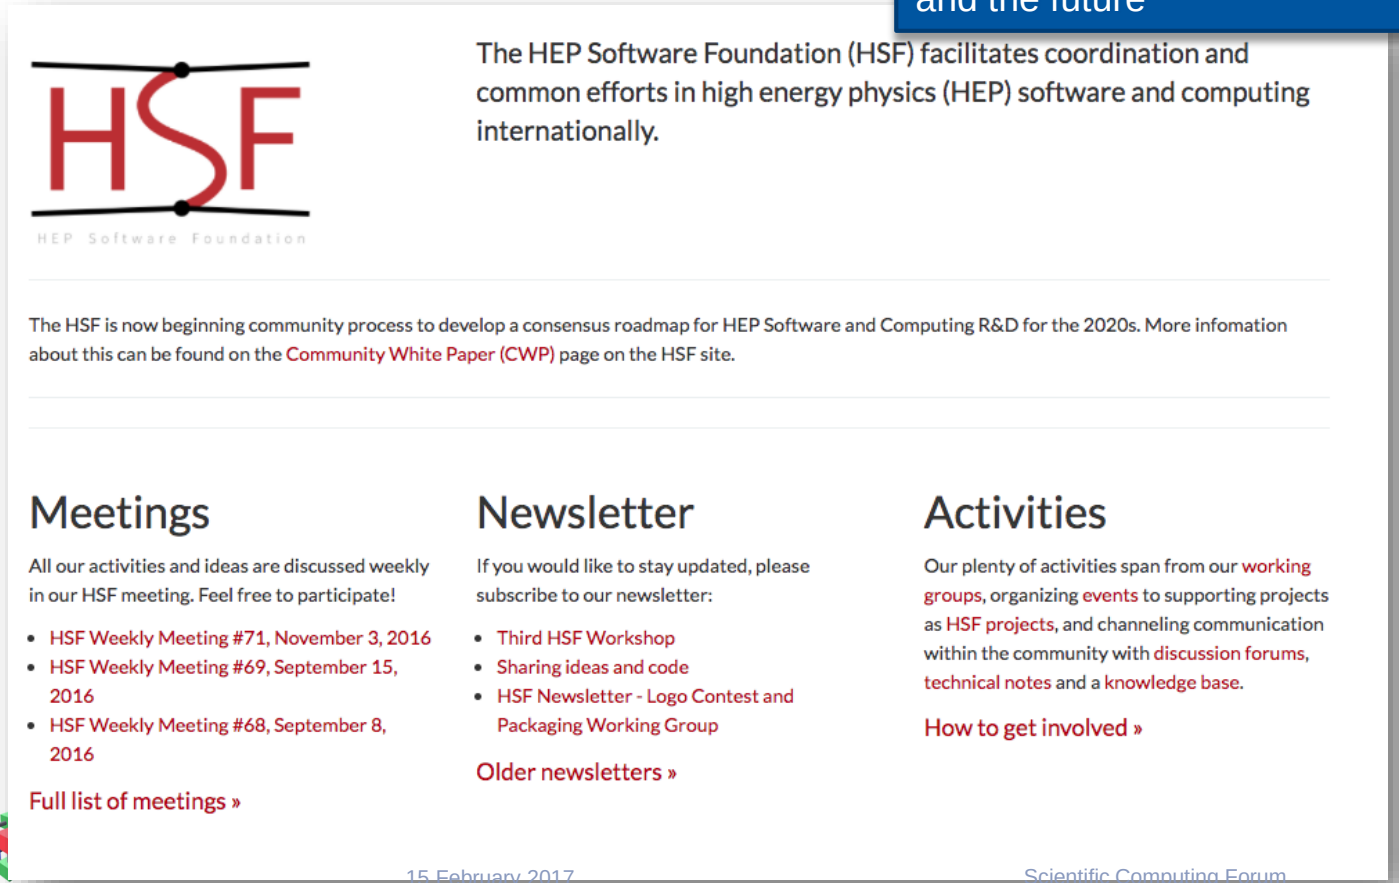

Software investment

- There are some (funded) activities in some countries
- Need to engage effort globally & coherently in HEP for this
  - Could use mechanism of original LCG project
    - Define project(s) with clear timescale and outcomes
    - Ask for contributions (of effort not cash)
- Use the HSF/CWP as the framework – we already have community engagement
- NB: Software and infrastructure must be treated together – the separation is a source of inefficiency

Steering/governance

Steering / governance

- Would need buy-in from all potential stakeholders
- Would like ~~oversight~~ steering from IT Heads of major HENP labs, projects, facilities, experiments
- Should have a globally recognized mandate
  - E.g. from ICFA (have had such before for grids); needed to get acceptance from some countries
- Governance should be very lightweight, and through consensus of the community, informing the steering group
- Role is to ensure that infrastructure and software:
  - Direction and evolution is suitable for the community;
  - Funding and effort – direct feedback to FA's or labs;
  - Can use HSF, HEPiX, specific projects, working groups.
- \+ AND: community-wide licensing, procurements, agreements
  - E.g. on commercial clouds – economy of scale
  - Address political issues – e.g. how to evolve funding models

{Experiment, project, facility} -specific

Experiment-specific

- Factor out the experiment- or facility- specific processes
  - E.g. The MoU and resource management (pledging) process of WLCG
    - Expect similar needs for other collaborations
  - New experiments will need to fund the computing resources they will require
- Note: this is not an attempt to use the same boxes for all experiments (like with EGEE, OSG, etc.)
  - Rather – try and use same infrastructure, tools, processes, software as far as possible so that new experiments are easier to support on existing facilities
  - Of course, this helps opportunistic use and sharing – but does not try and impose it

Practical steps

WLCG/HEP

- We could build a prototype of some of this now
  - HEP technology (EOS etc) – small demonstrators exist
    - EOS & dCache federations
- Opportunity of a new DC in Prévessin
- Software initiatives
  - HEP put HSF in place, recognising software as a crucial area
  - Must ensure that we invest sufficiently in people and skill building/retention

Prototyping

- Building a prototype data cloud/science cloud now would also allow us to provide additional value-add services to HEP
  - (Web-)FTS, cernbox, data archiving, DBOD, data preservation (Zenodo) and open access platform
    - All of these “as-a-service”
  - Could demonstrate to other sciences – to eventually contribute to a science cloud
    - And then include other collaborative services (Indico, Vidyo, etc.)
  - Leverages CERN’s & HEP key competencies and experience
- This would be a clear unique contribution of HEP to an EOSC infrastructure …