

Community White Paper: A Roadmap for HEP Software and Computing Status and Plans

Peter Elmer - Princeton University
(for the HSF Software Foundation)

11 October, 2016
CHEP 2016

A Software “Upgrade” for HL-LHC and 2020s HEP?

Looking forward to the next 10 years, we see a number of challenges for HEP software and computing:

- **Scale:** The HL-LHC will integrate 100 times the current data, with significantly increased data (pileup) and detector complexity.
- **Performance/cost:** Estimates of computing needs run faster than Moore’s Law by factors of 3-30
- **Technology/Market evolution:** the return of heterogeneity; technology change will also make it challenging to exploit Moore’s Law without software evolution.
- **Sustainability:** Most of the current software, which defines our capabilities, was designed 15-20 years ago: there are many software sustainability challenges.

Why Software? Software is *the* Cyberinfrastructure




Computer hardware is a consumable.
Software is what we keep, and invest in, over time.

Estimates of Resource Needs for HL-LHC (WLCG)



Data:

- Raw 2016: 50 PB → 2027: 600 PB
- Derived (1 copy): 2016: 80 PB → 2027: 900 PB



CPU:


- x60 from 2016

Technology at ~20%/year will bring x6-10 in 10-11 years

- ❑ Simple model based on today's computing models, but with expected HL-LHC operating parameters (pile-up, trigger rates, etc.)
- ❑ At least x10 above what is realistic to expect from technology with reasonably constant cost

(Slide from WLCG Workshop Intro, Ian Bird, 8 Oct, 2016)

Processor evolution and software impact



Clock Frequency vs Time

- Single core performance has stalled, leading to multi/manycore and specialization
- To even realize Moore's Law gains, we are pushed towards parallelization of algorithms and design for performance.
- The software designs and implementations themselves need to evolve, not just be recompiled

Back to heterogeneous systems?

Building the worldwide distributed LHC computing grid was largely made possible by the convergence on Linux on (commodity) Intel x86 processors around the year 2000. Building the WLCG at this scale in the heterogeneous workstation era would have been quite difficult. For better or for worse, heterogeneity is returning:

- Diversity of computing processor architectures (general purpose cores vs specialized processors)
- Owned vs commercial/cloud providers
- Some pressure to use systems traditionally designed for other types of applications (e.g. HPC/supercomputer as opposed to HTC/high-throughput systems)
- Possible further commoditizing market pressures (e.g. mobile)

What is software sustainability?

- **Dependent Infrastructure:** Will the infrastructure element continue to provide the same functionality in the future, even when the other parts of the infrastructure on which the element relies change?
- **Collaborative Infrastructure** Can the element be combined with other elements to meet user needs, as both the collaborative elements and the individual elements change?
- **New Users:** Is the functionality and usability of the infrastructure element clearly explained to new users? Do users have a mechanism to ask questions and to learn about the element?
- **Existing Users:** Does the infrastructure element provide the functionality that current users want? Is it modular and adaptable so that it can meet the future needs of the users?
- **Science:** Does it incorporate and implement new science and theory as they develop?

HEP Software Foundation (HSF)

The HSF (<http://hepsoftwarefoundation.org>) was created 1.5 years ago as a means for organizing our community to address the software challenges of future projects such as the HL-LHC. The HSF has the following objectives:

- Catalyze new common projects
- Promote commonality and collaboration in new developments to make the most of limited resources
- Provide a framework for attracting effort and support to S&C common projects (new resources!)
- Provide a structure to set priorities and goals for the work



HEP Software Ecosystem



IgProf

FroNTier

FairRoot



FastJet

RooStats

EvtGen



Examples, definitely incomplete!



Recent/Nascent Cross-experiment Collaborations

- Experiment frameworks
 - Gaudi, FAIRRoot, CMSSW/Art
- Common Conditions Data Project
 - Discussion/cooperation between ATLAS, Belle II, CMS and LHCb
- Common Software Build and Packaging Tools efforts
 - Working group of HSF comparing HEP and non-HEP solutions
- Cooperation on Reconstruction Software
 - “Connecting the Dots” tracking workshop, HSF sessions
- AIDA2020 (EU funded)
 - DD4hep for detector description, PODIO data model library (LCD, FCC, potentially LHCb)
- DIANA (Data Intensive ANALysis) (NSF Funded)
 - 4-year project on analysis software, including ROOT and its ecosystem

Defining Longer-term Strategy

- HL-LHC computing requires a major 'software upgrade'
- A Community White Paper (CWP) on the overall strategy and roadmap for software and computing has been proposed
 - Initiated as WLCG charge to the LHC experiments and HSF as a step towards the LHC experiment TDRs in advance of HL-LHC
 - The scope should not be restricted only to HL-LHC
 - Some early software components could be built, tested and used by experiments in LHC Run3
- Organised by the HEP Software Foundation (HSF)
- Paper to be delivered by Summer 2017
- It should play a role in discussing possible funding scenarios for a "software upgrade".

Community White Paper (CWP)

- The CWP should identify and prioritise the software research and development investments required:
 - to achieve improvements in software efficiency, scalability and performance and to make use of the advances in CPU, storage and network technologies
 - to enable new approaches to computing and software that could radically extend the physics reach of the detectors
 - to ensure the long term sustainability of the software through the lifetime of the HL-LHC
- We need to engage the HEP community in this process through a series of workshops
 - Initiated as an HL-LHC planning process
 - Aiming for a broader participation (LHC, neutrino program, Belle II, linear collider so far)

Likely constraints to fund a “Software Upgrade”

It appears unlikely that significant increases in investments in software will be made by funding agencies purely from particle physics budgets and/or into individual experiments. Other opportunities do perhaps exist, but often imply constraints, for example:

- Investments into software impacting multiple experiments
- Investments into development with impact beyond particle physics
- Investments into development permitting use of computing facilities (e.g. HPC) planned for other non-HEP purposes
- Investments requiring collaborations with Computer Science or Industry

Building the LHC software in use today was possible without too many such constraints. The good news is that the community (with an existing LHC computing system) is better positioned today to make effective progress even with such constraints.

Status

The proposal for a general Community Roadmap has been widely discussed with all of the LHC experiments and the HEP Software Foundation (HSF). There is broad support for the idea.

The CWP roadmap plan, to be carried out by HSF, was presented to the LHCC. It fits with the current notion of HL-LHC computing TDRs in ~2019-2020.

WLCG has produced a charge for this CWP to the HSF and the LHC experiments (see separate link) with an aim to complete it by the end of August, 2017.

The HSF has begun the process of organizing working groups, engaging HEP beyond the LHC experiments and planning for dedicated workshops. Sessions at existing meetings can also be used when possible.

Community Roadmap Process

We propose a series of workshops over the next year to build the community roadmap:

- Initial presentation and organization this month (at WLCG workshop and CHEP, etc.)
 - Flesh out the charges and attract interested individuals to the WGs
- A “kick-off” workshop at UC San Diego on 23-26 Jan 2017
 - Start real work after a few months post-CHEP gestation in the WGs
 - Discussions on more controversial topics, find path to consensus
 - Develop plans and responsibilities for delivering white paper by summer 2017
- Possible “topical” workshops between Jan-Jun 2017, building on existing community activities when possible (e.g. DPHEP, Reco Algorithms Forum/CTD, IML)
- A final workshop in summer 2017 (in Europe, near CERN?)

What should the community roadmap process accomplish?

Going back to the subset of HSF goals I listed earlier:

- Catalyze new common projects
- Promote commonality and collaboration in new developments to make the most of limited resources
- Provide a framework for attracting effort and support to S&C common projects (new resources!)
- Provide a structure to set priorities and goals for the work

The workshop process, the community roadmap white paper and (simultaneously) the pursuit of specific plans/proposals will support precisely these goals.

Possible routes to a “Software Upgrade”

- If we are aiming at a larger “software upgrade” project towards the HL-LHC, an additional ingredient is to find (or liberate/reallocate) the resources to realize this roadmap.
- We need both initial exploratory R&D and eventual development projects!
- In the US, both the NSF and the DOE have at least the notion of eventual resources and/or organization for new common projects in HEP (NSF: SI2, DOE: HEP CCE)
 - The US NSF has funded a “conceptualization” (planning) project with a possible path towards a “Software Institute”.
 - The US DOE has seeded the “Center for Computing Excellence” with some initial resources.
- We hope that a clear community roadmap will bring these and other partners together for an HL-LHC software upgrade.

Practicalities: HSF Google Groups

The following Google Groups are relevant:

- Group for discussion of Community White Paper
 - <https://groups.google.com/forum/#!forum/hsf-community-white-paper>
- General announcement group for community messages (low traffic)
 - <https://groups.google.com/forum/#!forum/hep-sw-comp>
- Community Discussion list
 - <https://groups.google.com/forum/#!forum/hep-sf-forum>
- Specific group for US NSF Software Institute Conceptualization
 - <https://groups.google.com/forum/#!forum/s2i2-hep>

Practicalities: Charges for CWP Working Groups

- Over the next weeks we will be formulating charges for the CWP working groups.
- Templates for drafting these charges are in google docs.
- The overall WLCG HSF charge and links to individual WG charge google docs (in preparation) can be found at:
 - <http://bit.ly/2dcZZqa>
- To view and/edit these charges you will need to be subscribed to:
 - <https://groups.google.com/forum/#!forum/hsf-community-white-paper>
- These google docs are just to draft the charges and allow people to self-organize and start discussions. Eventual proper documents (e.g. in latex) can switch elsewhere (e.g. github).

Practicalities: Possible Working Groups

Detector Simulation	full and fast simulations, hi-pileup environments
Triggering	algorithms, GPUs and/or FPGAs
Event Reconstruction	new approaches to event reconstruction
Visualization	tools for data analysis, education, and outreach
Data Access and Management	scaling to the exabyte level
Workflow and Resource Management	millions of jobs in heterogenous systems
Physics generators	better models, better precision, code optimisations
Data Analysis and Interpretation	efficient use of many-core, modern techniques
Data and Software Preservation	preservation and reuse of data and software
Software Development, Deployment and Validation/Verification	improved modularity and quality, contribution
Computing Models, Facilities, Distributed Computing Various Aspects of Technical Evolution (Software Tools, Hardware)	range of possible models, costing, technology
Security and Access Control	
Careers, Staffing and Training	perhaps in a separate concurrent white paper
Machine Learning	
Conditions Database	
Event Processing Frameworks	

More details in links at <http://bit.ly/2dcZZqa>

Practicalities: Documents

- The end goal here is a single CWP roadmap for the community.
- The first step is building working groups and defining specific charges
- The process of putting together the CWP should generate a series of narrower topical documents from the working groups. (Much like the Snowmass process, for example.)
- Existing public documents are something we will build upon, e.g. the Snowmass Computing documents, the DOE HEP-CCE documents, the WLCG Run2 Computing Model Update, the CERN Openlab whitepaper, etc.
- The specific mechanics of assembling the CWP document itself will be defined at the Jan2017 SDSC HSF CWP workshop

Discussion questions

- How can we best organize to produce a consensus roadmap for the Community White Paper? What is missing from the proposed process?
- How can we best address the three CWP goals?
 - to achieve improvements in software efficiency, scalability and performance and to make use of the advances in CPU, storage and network technologies
 - to enable new approaches to computing and software that could radically extend the physics reach of the detectors
 - to ensure the long term sustainability of the software through the lifetime of the HL-LHC
- In practice how do we re-examine the organizational processes by which the HEP community and the experiments collaborate?
- What opportunities do we have for funding a “software upgrade” to address these challenges?

HEP Software Ecosystem



IgProf

FroNTier

FairRoot



FastJet

RooStats

EvtGen



Examples, definitely incomplete!

