The DIANA/HEP and S2I2-HEP projects, the HEP Software Foundation and the Community White Paper Towards Sustainable Software

Peter Elmer - Princeton University

13 January, 2017
INFN-Padova Discussion

In this presentation I will talk about several things:

- Software Challenges and Sustainability for HEP in the 2020s
- The DIANA/HEP project
- The S2I2-HEP "Software Institute" conceptualization project
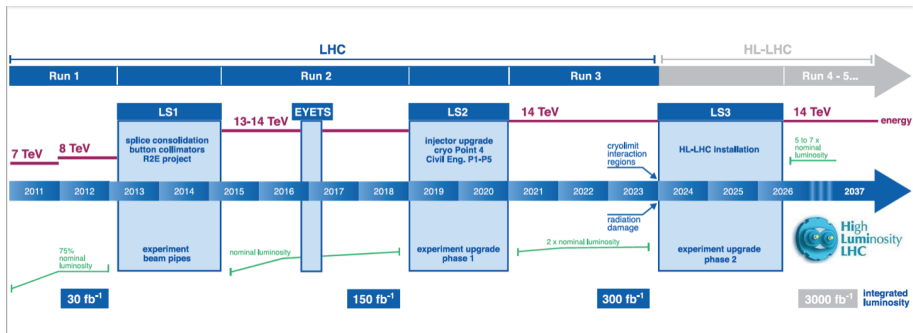- The HEP Software Foundation and the "Community White Paper" roadmap project

I'll also explain the ways in which these are related...

CERN Accelerator Timeline
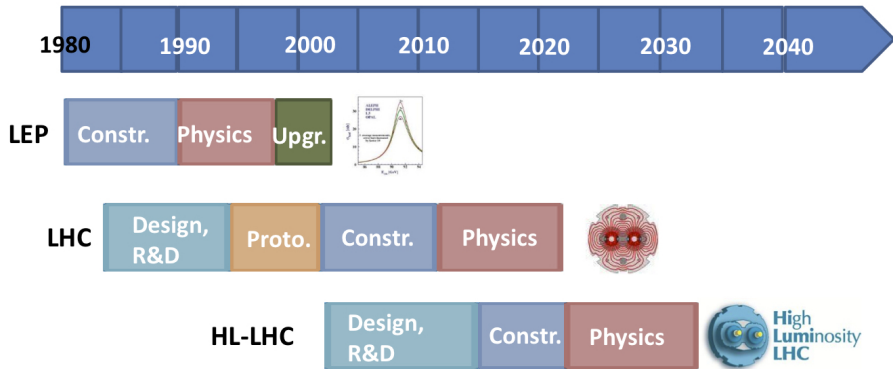
Various concepts also exist for subsequent machines.

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.

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

Data:
- Raw 2016: 50 PB → 2027: 600 PB
- Derived (1 copy): 2016: 80 PB → 2027: 900 PB
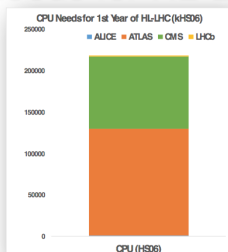
CPU:
- x60 from 2016

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
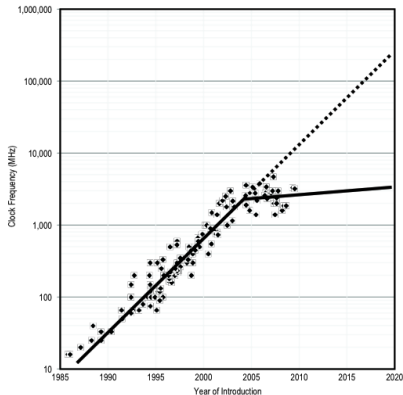
Clock Frequency vs Time

- 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?

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.
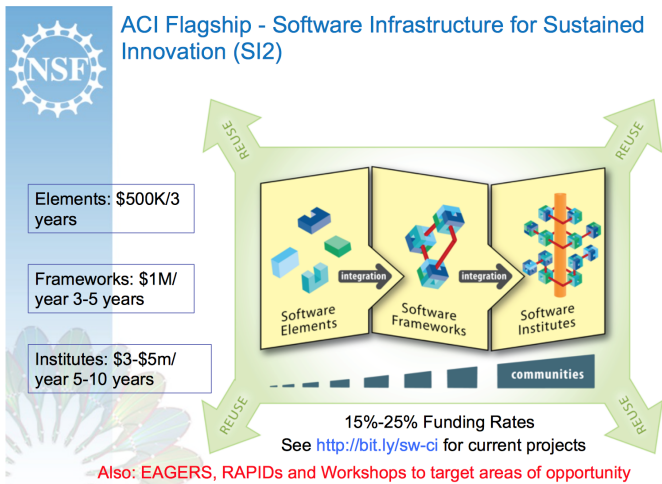
Examples, definitely incomplete!

Plus 15-20M Source Lines of Code (SLOC) of "experiment specific" codes, as well as dependencies on non-HEP scientific software.

Software Infrastructure for Sustained Innovation (SI2)

ACI Flagship - Software Infrastructure for Sustained Innovation (SI2)

Elements: $500K/3 years

Frameworks: $1M/ year 3-5 years

Institutes: $3-$5m/ year 5-10 years

Software Elements

Software Frameworks

Software Institutes

communities

15%-25% Funding Rates
See http://bit.ly/sw-ci for current projects

Also: EAGERS, RAPIDs and Workshops to target areas of opportunity

DIANA/HEP is a "Software Framework" (SSI). The S2I2-HEP project is a planning (conceptualization) "Software Institute" award.

NSF SI2 Award Classes

The NSF SI2 program includes three classes of awards:

- **Software Elements (SSE)** target small groups that will create and deploy robust software elements for which there is a demonstrated need that will advance one or more significant areas of science and engineering.

- **Software Frameworks (SSI)** target larger, *interdisciplinary* teams organized around the development and application of elements of common software infrastructure aimed at solving common research problems. These awards will result in sustainable community software frameworks serving a diverse community.

- **Scientific Software Innovation Institutes (S^2I^2)** focus on the establishment of long-term hubs of excellence in software infrastructure and technologies, including elements and frameworks, that will serve a *research community of substantial size and disciplinary breadth*.

The DIANA/HEP project

- Data Intensive ANAlysis for High Energy Physics (DIANA/HEP)
- The primary goal of DIANA/HEP is to develop state-of-the-art tools for experiments which acquire, reduce, and analyze petabytes of data.
- DIANA is not a piece of software itself, but a collaborative project to improve and extend analysis tools as sustainable infrastructure for the community.
- Funded by NSF "Software Infrastructure for Sustained Innovation" (SI2) program
- 4-year project, 6-7FTE total
- Princeton, NYU, UCincinnati, U.Nebraska-Lincoln
- The PIs are involved in Atlas, CMS and LHCb

The DIANA/HEP Project (http://diana-hep.org)

The DIANA/HEP project focuses on improving performance, interoperability, and collaborative tools through modifications and additions to ROOT and other packages broadly used by the HEP community.

Website: `http://diana-hep.org`

Google group: `https://groups.google.com/forum/#!forum/diana-hep`

Github: `https://github.com/diana-hep`

DIANA/HEP team

DIANA Team

Project Team

- Peter Elmer (Lead PI) - Princeton University, Department of Physics
- Brian P. Bockelman (PI) - University of Nebraska-Lincoln, Department of Computer Science and Engineering
- Kyle Cranmer (PI) - New York University, Department of Physics & Center for Data Science
- Michael D. Sokoloff (PI) - University of Cincinnati, Department of Physics
- Jinyang Li (Senior Personnel) - New York University, Computer Science Department
- David Lange - Princeton University, Department of Physics
- Gilles Louppe - New York University, Department of Physics & Center for Data Science
- James Pivarski - Princeton University, Department of Physics
- Eduardo Rodrigues - University of Cincinnati, Department of Physics
- David Abdurachmanov - University of Nebraska-Lincoln, Department of Computer Science and Engineering
- Zhe Zhang - University of Nebraska-Lincoln, Department of Computer Science and Engineering (Ph.D. Student)
- Chien-Chin Huang - New York University, Computer Science Department (Ph.D. Student)
- Lukas Heinrich - New York University, Department of Physics (Ph.D. Student)

Associated team members

- Vassil Vassilev - Princeton University & ROOT Intel Parallel Computing Center

DIANA/HEP collaborators

Collaborators

- The ROOT team at CERN and Fermilab
- CMS Big Data Project: Oliver Gutsche, Matteo Cremonesi, Nhan Tran, Jim Kowalkowski, and Saba Sehrish - Fermilab
- Histogrammar: Alexey Svyatkovskiy - Princeton University
- Scikit-HEP: Noel Dawe - University of Melbourne, Vanya Belyaev - ITEP, and Sasha Mazurov - University of Birmingham
- Spark-ROOT: Viktor Khristenko - University of Iowa
- Scope-aaS: Jin Chang and Igor Mandrichenko - Fermilab
- Scope-GPU: Roger Rusack and Peter Hansen - University of Minnesota
- Carl: Juan Pavez, Cyril Becot - New York University, Lukas Heinrich - New York University
- Scikit-Optimize: Manoj Kumar - New York University, Tim Head, Noel Dawe - University of Melbourne

Example DIANA talks/activities

Activities and Research Products

Software products

- DIANA/HEP organization on GitHub
- carl docs `DOI` `10.5281/zenodo.47798`
- histogrammar - web page, docs

Papers

- Carl in the Journal of Open Source Software `JOSS` `10.21105/joss.00011`

DIANA Topical Meetings

The DIANA team organizes periodic meetings on various topics of interest, with both invited, externals speakers and DIANA team contributions. These and other DIANA meetings can be found in the INDICO agenda system.

Presentations by the DIANA team

- 14 Dec 2016 - Data Plumbing (whiteboard talk), Jim Pivarski, LPC Coffee Hour
- 5 Dec 2016 - Focus groups: physicist attitudes toward software, Jim Pivarski, DIANA Meeting
- 14 Nov 2016 - TTreeProcessor: A toy framework for parallel ntuple processing, Brian Bockelman, DIANA Meeting
- 7 Nov 2016 - Overview of file formats, Jim Pivarski, DIANA Meeting
- 7 Nov 2016 - Reading ROOT data in Java and Spark, Jim Pivarski, ROOT Team meeting; follow up with CERN IT
- 30 Oct 2016 - Plotting data on GPUs with Histogrammar, Jim Pivarski, Parallel Kalman Meeting
- 10 Oct 2016 - Exploring Compression Techniques for ROOT IO, Zhe Zhang CHEP 2016
- 20 Oct 2016 - Histogrammar: Making Histograms Functional, Jim Pivarski Chicago Hadoop User's Group (CHUG)
- 21 Sep 2016 - Reconstructing Particle Trajectories in High Energy Physics with Xeon and Xeon Phi, David Abdurachmanov, Intel Xeon Phi User Group Meeting 2016
- 16 Sep 2016 - Expressing Complex Data Aggregations With Histogrammar, [talk abstract, talk video], Jim Pivarski, Strange Loop Conference
- 18 Aug 2016 - Computing for Data Analysis, Jim Pivarski, LPC Physics Forum
- 18 Jul 2016 - Potential Issues with Julia, Jim Pivarski, DIANA Meeting
- 20 Jun 2016 - High-level analysis scripts with low-level performance, Jim Pivarski, DIANA Meeting

DIANA Topical Meetings (https://indico.cern.ch/category/7192/)

December 2016

| | 12 Dec | First Explorations NEW |
| | 05 Dec | Software Development, Sustainability and Evolution |

November 2016

| | 14 Nov | Toward Vectorized, Parallel NTuple processing |
| | 07 Nov | Avro and Numpy for Spark and machine learning |

October 2016

| | 31 Oct | DIANA Meeting - Event generation and fitting on GPUs |

September 2016

| | 19 Sep | DIANA Meeting - Python/ROOT interoperability |

July 2016

| | 25 Jul | DIANA Meeting - Focus groups, surveys and building a roadmap |
| | 18 Jul | DIANA Meeting - The Julia Language |

June 2016

| | 20 Jun | DIANA Meeting - Analysis script language |

Upcoming DIANA topical meetings

February 2017

	20 Feb	DIANA Meeting - Scikit-HEP
	13 Feb	DIANA Meeting - Pachyderm and Go
	06 Feb	DIANA Meeting - Scikit-Optimize

January 2017

| | 30 Jan | DIANA Meeting - Machine learning with TMVA and Scikit-Learn |

NSF SI2-S2I2 Software Institute

NSF SI2-S2I2 includes two subclasses of awards:

- Conceptualization Awards - which are planning awards aimed at organizing an interdisciplinary community and understanding their software requirements and challenges ($500k, 1-2 years)

- Implementation Awards - which will be made to implement community activities that support software infrastructure, for example, such as those developed by the conceptualization awards ($3-5M/year, 5 years)

- NSF funded two software institute implementations last summer (2016):

- https://www.nsf.gov/news/news_summ.jsp?cntn_id=189347

- And they also funded our S^2I^2 conceptualization project:
  - Conceptualization of an S^2I^2 Institute for High Energy Physics
  - http://cern.ch/elmer/s2i2-2015-nsf-proposal.pdf

S2I2-HEP

The primary goal of the S2I2-HEP conceptualization project (`http://s2i2-hep.org`) is to produce a well-defined strategy for developing the software and computing models for use in high energy physics (HEP), in particular for the experiments collecting the very large data sets anticipated in the "High-Luminosity Large Hadron Collider" (HL-LHC) era of the 2020s.

Specifically the S2I2-HEP project will identify potential areas where U.S. university personnel can lead in key areas of software development to help realize the full potential of the HL-LHC program.

However HEP and the LHC are global projects, so no long-term planning exercise can exist in isolation, thus we are also pursuing a wider HEP community roadmap for software and computing in the 2020s.

S^2I^2 Conceptualization Awards

S^2I^2 Conceptualization Awards are *planning* awards aimed at organizing an interdisciplinary community and understanding their software requirements and challenges. Example activities that may be undertaken as part of this award include focused workshops, special sessions at professional meetings, sandpits, focus groups, etc.

The product of a conceptualization award will be a **Strategic Plan** for enabling science and education through a sustained software infrastructure that will be freely available to the community.

The Strategic Plan resulting from the conceptualization phase is expected to serve as the conceptual design upon which a subsequent S^2I^2 Implementation proposal could be based.

Because an NSF effort cannot stand in isolation to DOE and international efforts, the process we proposed delivers also a broader **Community White Paper**.

HEP Software Foundation (HSF)

The HSF (http://hepsoftwarefoundation.org) was created in early 2015 as a means for organizing our community to address the software challenges of future projects such as the HL-HLC. 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

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

Community White Paper (CWP)

- The CWP will 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

- The HSF is engaging the HEP community to produce the CWP via a "community process"
  - Initiated as an HL-LHC planning process
  - Aiming for a broader participation (LHC, neutrino program, Belle II, linear collider so far)

Detector Simulation, Triggering, Event Reconstruction and Visualization

Challenges surrounding high pile-up simulation, including the CPU resources needed for large statistics samples needed to compare with data from high trigger rates, high memory utilization, generation and handling of the large (min-bias) samples needed to achieve accurate description of high pile-up collision events, and a flexible simulation strategy capable of a broad spectrum of precision in the detector response, from "fast" (e.g. parametric) simulation optimized for speed to full simulation in support of precision measurements and new physics searches (e.g. in subtle effects on event kinematics due to the presence of virtual particles at high scale). Software required to emulate upgraded detectors (including the trigger system) and support determination of their optimal configuration and calibration. • Software in support of triggering during the HL-LHC, including algorithms for the High-level Trigger, online tracking using GPUs and/or FPGAs, trigger steering, event building, data "parking" (for offline trigger decision), and data flow control systems. • New approaches to event reconstruction, in which the processing time depends sensitively on instantaneous luminosity, including advanced algorithms, vectorization, and execution concurrency and frameworks that exploit many-core architectures. In particular, charged particle tracking is expected to dominate the event processing time under high pile-up conditions. • Visualization tools, not only in support of upgrade detector configurations and event displays, but also as a research tool for data analysis, education, and outreach using modern tools and technologies for 3D rendering, data and geometry description and cloud environments.

Data Access and Management, Workflow and Resource Management

Data handling systems that scale to the Exabyte level during the HL-LHC era and satisfy the needs of physicists in terms of metadata and data access, distribution, and replication. Increasing availability of very high speed networks removes the need for CPU and data co-location and allows for more extensive use of data access over the wide-area network (WAN), providing failover capabilities, global data namespaces, and caching. • Event-based data streaming as complementary to the more traditional dataset-based or file-based data access, which is particularly important for utilizing opportunistic cycles on HPCs, cloud resources, and campus clusters where job eviction is frequent and stochastic. • Workflow management systems capable of handling millions of jobs running on a large number of heterogeneous, distributed computing resources, with capabilities including whole-node scheduling, checkpointing, job rebrokering, and volunteer computing. • Systems for measurement and monitoring of the networking bandwidth and latency between resource targets and the use of this information in job brokering. • Software-defined networking technologies which enable networks to be configurable and schedulable resources for use in the movement of data.

Physics generators, Data Analysis and Interpretation, Data and Software Preservation

There are many theory challenges in the HL-LHC era, among them are improving the precision of SM calculations, better estimation of systematic uncertainties, and elucidation of promising new physics signals for the experiments. Software needed to make connection between observations and theory include matrix element generators, calculation of higher-order QCD corrections, electroweak corrections, parton shower modeling, parton matching schemes, and soft gluon resummation methods. Physics generators that employ concurrency and exploit many-core architectures will play an important role in HL-LHC, as well better sharing of code and processing between LHC experimenters and phenomenologists. • Data analysis frameworks that include parallelization, optimized event I/O, data caching, and WAN-based data access. Analysis software that employs advanced algorithms and efficiently utilizes many-core architectures. • Tools and technologies for preservation and reuse of data and software, preservation and re-interpretation of physics results, analysis providence and workflow ontologies, analysis capture, and application packaging for platform abstraction. • Future software repositories and build platforms that leverage advances in these areas and improved software modularity and quality control that will allow a broader community of people to effectively contribute to software in the HL-LHC era.

Practicalities: CWP Process

- The end goal here is a single (consensus) CWP roadmap for the community.
- Finding consensus in a large community is a difficult task: broad participation and visibility/transparency are key elements
- The process being used largely mirrors that used in the "decadal survey" process in high energy physics
- Working groups self-organize, with encouragement from institutions and projects/experiments/etc.
- A series of workshops is planned over about 9 months to allow topics to be explored, sometimes overloading CWP discussions onto preexisting meetingd
- Contributions along the way can come in the form of "white papers" by individuals/groups/projects/institutions
- Based on the ideas emerging from the discusions, workshops and white papers, a consensus CWP document will be written.

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: 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	range of possible models, costing
Various Aspects of Technical Evolution (Software Tools, Hardware)	
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://hepsoftwarefoundation.org/cwp.html

HSF Community White Paper Workshop at SDSC

- The kick-off workshop for the HSF CWP process will be on 23-26 January 2017 at SDSC/UCSD
- http://indico.cern.ch/event/570249/
- People from many HEP experiments (LHC and beyond)
- A key element in the "community process" to form a consensus
- We are looking for opportunities to introduce new ideas into the HEP discussions
- An early strawman agenda is posted (next slides, but evolving quickly so online version has the latest)
- Registration is open (see URL above)

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" (http://s2i2-hep.org)
    - The US DOE has seeded the "Center for Computing Excellence" with some initial resources. (http://hepcce.org)
- We hope that a clear community roadmap will bring these and other partners together for an HL-LHC software upgrade.
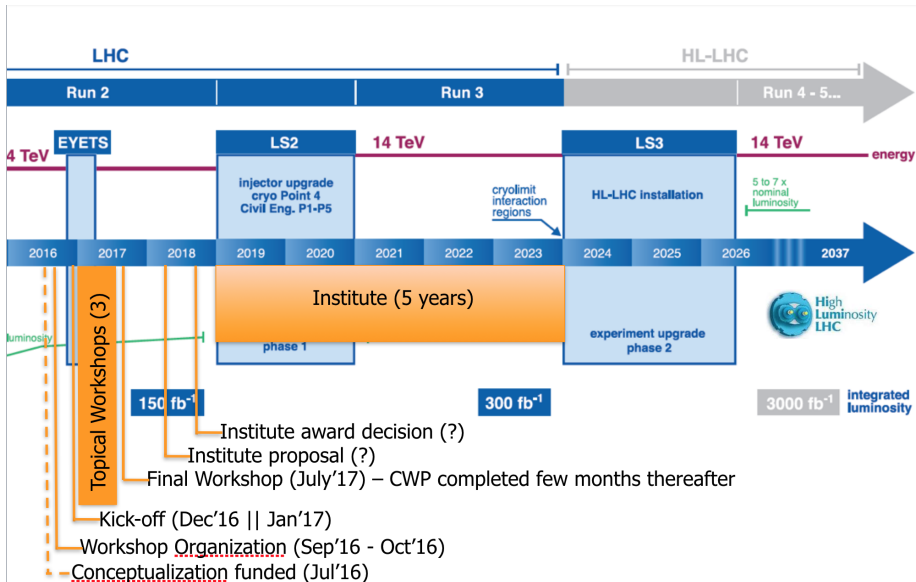
Monday Plenary Agenda (evolving)

Mon 23/01	Tue 24/01	Wed 25/01	Thu 26/01	All days

Print   PDF   Full screen   **Detailed view**   Filter
Session legend

◯ Computing Models, F...   ● Introduction   ◯ Introduction and Over...   ◯ Invited Talks   ✕

08:00

Workshop Welcome	
	08:30 - 08:35
Status of HSF and Community White Paper Plans	*Dr. Peter Elmer*
	08:35 - 09:05
LHC Future Challenges - Software/Computing	*Ian Bird*
	09:05 - 09:30
LHC Future Challenges - Physics	
	09:30 - 10:00
Challenges for the Fermilab neutrino and muon programs	*Robert Kutschke*
	10:00 - 10:15
Future software challenges at Jefferson Lab	*Amber Boehnlein*
	10:15 - 10:30

09:00

10:00

11:00

HEP experiments and software status in China	*Prof. Wenjing Wu*
	11:00 - 11:15
Software plans and challenges for Linear Collider	*Frank-Dieter Gaede*
	11:15 - 11:25
NSF View	*Dr. Bogdan Mihaila*

Timetable

Mon 23/01	Tue 24/01	**Wed 25/01**	Thu 26/01	All days

Print   PDF   Full screen   **Detailed view**   Filter

Session legend

○ Data Access and Mana...   ● Data Analysis and Inter...   ○ Data Analytics and Mac...   ○ Data Analytics and Mac...   ✕

see more...

08:00

Panel: Data Centres and Facilities Technologies (Clouds, Networks, Storage)

09:00

08:30 - 09:30

Data Analytics and Machine Learning	Data Access and Management	Visualization	Parallel 12

10:00

11:00

12:00

Practicalities: Contributing White Papers

- The end goal here is a single CWP roadmap for the community.
- Along the way people are encouraged to put ideas on the table by submitting topical white papers.
- http://hepsoftwarefoundation.org/cwp-whitepapers.html
- Note in particular the request for contributed white papers on "Computing Models, Facilities and Distributed Computing" by 15 January, 2017
- http://hepsoftwarefoundation.org/cwp/ CWPWhitePaperSolicitation.pdf
- Existing public documents are also 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. Feel free to send links to these, too.

Summary

- We have a significant investment in software, it embodies the core of our intellectual property and the real cyberinfrastructure.
- Significant challenges of scale, performance, technology and long term sustainability exist as we face the projects of the 2020s.
- DIANA/HEP is a pilot project to understand how to build sustainable software and software communities
- The LHC and wider HEP communities are also executing a planning process for medium/long term software and computing R&D
- The process should produce the "Community White Paper", a roadmap for HEP Software and Computing for the 2020s
- In parallel we would like to investigate funding and collaboration possibilities for eventual projects and activities
- Please consider participating in the CWP process, the WGs and the workshops (beginning with the HSF workshop SDSC/UCSD in January 2017) and to put forward ideas in white papers