Clarification of types of computing staff for experimental particle physics computing infrastructure – and historical support routes

Prepared in consultation with GridPP, ATLAS, CMS, LHCb

Software infrastructure is required by particle physics experiments, just as is the physical computing infrastructure, to acquire, handle, process and analyse data. This short document is intended to clarify the different roles for which computing staff are required.

The table-diagram below shows these different roles. It shows the difference between those which are only supportable via PPGP (and PPRP) and those which are supported through GridPP.

This table emphasises the distinction between the different types of work done in each layer, but there is of course close communication between the people involved in experiment-specific effort and WLCG, especially between E3 and W1. Individuals work in several layers.
[Note: this table does not include local institute computing support posts]

Layer		Experiment – 1	Experiment – 2	Experiment - 3	
E1	Consolidated Grants or Project Grants	Physicists using analysis frameworks and adding analysis specific code	Physicists using analysis frameworks and adding analysis specific code	Physicists using analysis frameworks and adding analysis specific code	Vertical Layer Research software Engineering ⇒ Development and upgrading ↵ Feeds back
E2		Experiment specific reconstruction and analysis software	Experiment specific reconstruction and analysis software	Experiment specific reconstruction and analysis software	
E3		Experiment specific effort for production computing, data management, trigger and online computing.	Experiment specific effort for production computing, data management, trigger and online computing.	Experiment specific effort for production computing, data management, trigger and online computing.	
	WLCG (Worldwide LHC Computing Grid)				
W1	GridPP	Common distributed computing software infrastructure deployment on WLCG, operations, support and development. Global services such as security response, service registries, monitoring and accounting. Software verification and rollout.			
W2		Physical infrastructure deployment and operations staff			

Layers E1, E2, E3 lie entirely within each of the Experiments themselves

- Layer E1 captures the software effort by physicists working on analysis leading to the published outputs. This is shown for completeness, although is not normally labelled as software effort.

- Layer E2 includes all of the reconstruction and analysis software and tools required to manipulate and analyse the data, and typically includes:
    - Detector calibration and alignment;
    - data quality monitoring and assurance;
    - data reconstruction and streaming;
    - Monte Carlo Simulation software (MC);
    - trigger algorithms;
    - luminosity and non-collision background software;
    - event visualisation tools.

  This work is typically carried out by compute-experienced physicists who are expert in relevant areas – e.g. track reconstruction.

- Layer E3 includes all of the online and offline computing infrastructure required for an experiment to take data, to run and manage the payloads from layers E1 and E2, and typically includes:
    - Online computing for data taking;
    - the trigger framework;
    - experiment reconstruction and analysis computing framework;
    - workload management system;
    - bulk experiment wide data management (storage, transfer, replication, cataloguing);
    - production computing operations (running many thousands of simultaneous jobs);
    - testing and monitoring frameworks;
    - computing resource and model evolution;
    - DMP (data management plan) enactment: data preservation, open data compliance
    - analysis curation and preservation;
    - information protection, software training & documentation;
    - operation shifts for data processing and quality monitoring (fair share expected).

  This work is typically carried out by dedicated experiment offline and online computing teams (with ever more convergence of the two)

E1-E3 are experiment specific. E1-E3 are typically requested either through Consolidated Grants (PPGP) for exploitation phase or the PPRP for Construction Projects.

E1-E3 cannot be requested as part of GridPP, as it is experiment specific work.

Layers W1 and W2 lie entirely within the distributed computing infrastructure known as the WLCG (Worldwide LHC Computing Grid)[1]

[1] The WLCG is the largest distributed scientific computing infrastructure in the world. It is also one of the largest distributed systems per se (outside enterprises such as Google, social media). Many of the functions listed are formal processes required not only to function on such a scale, but to ensure regulatory compliance.

- Layer W1 provides the services needed to run experiment workloads. These include:
    - The WLCG fabric (Compute Elements, Storage Elements, ..);
    - (high performance) data transfer services over wide area network;
    - configuration management
    - experiment software distribution to the WLCG sites;
    - storage systems, data federation components;
    - resource monitoring and accounting and resource information publication;
    - security policy and incident response;
    - operations functions (QA, ticketing, incident escalation process, coordination);
    - virtual organisation management system (VOMS), certification authority;
    - software verification and rollout facilities;
    - enabling tools and applications (e.g. for benchmarking, repositories);
    - legal compliance with data protection laws
    - installation of experiment-specific services at sites
    - response to site-specific problem tickets raised by those in E3;
    - training.

- Layer W2 is required to deploy and operate the CPU, disk and storage at the WLCG sites including:
    - Selection, procurement and installation and commissioning of equipment;
    - hardware maintenance and support;
    - operating system deployment and patching;
    - database system management and operation;
    - storage platform management and operation
    - exception monitoring and response (operations and security);
    - networking;
    - machine room operations and maintenance.

These layers are, or will be, in common for ALICE, ATLAS, CMS, LHCb, Lux-Zeplin, T2K, NA62,SNO+,MICE, DUNE, HyperK, Calice, Comet, ILC/CLIC and any new experiments to come.

An important distinction is that all of W1 would still be needed in the extreme situation that all physical resources were provided by a third party (e.g. commercial cloud[2]). It is only W2 which is linked to the physical hardware itself.

W1 and W2 effort is typically requested through GridPP.

W1 and W2 have so far not been supported through PPGP or PPRP.

Vertical layer: Development and Upgrades (Research Software Engineering[3])

The vertical layer is software engineering to develop and evolve the software. This is not part of the running production and analysis pipeline, but is the equivalent of an “upgrade” to the computing

[2] More on this is not relevant to this document, but in summary use of commercial cloud is not possible at present due to both technical limitations and cost model limitations (much more expensive)

[3] This vertical layer is referred to as “Research Software Engineering” (RSE) to make manifest the professional nature of the work and its equivalence with the term as used by EPSRC and within the STFC DiRAC HPC context; the RSEs are usually embedded in the projects in order to be effective.

infrastructure, just as one upgrades detector hardware to meet higher luminosity conditions. Such "upgrading" is required to adapt to the data rates expected from Run-3 and beyond of the LHC, as detailed in the HEP Software Foundation roadmap (HSF Roadmap) https://arxiv.org/abs/1712.06982

Examples include:

- R&D on architectures, technologies, frameworks, networks and algorithms;
- parallelising frameworks;
- distributed computing middleware development;
- moving services from the experiment specific to the WLCG wide.
- optimising, refactoring and writing new code to cope with higher rates and new hardware/operating systems/architectures;
- developing fast simulation capability needed to reduce production time;
- optimising and refactoring physics models in common simulation tools;
- event model and workflow evolution to reduce data storage requirements;
- developing physics generators to match experimental precision, and increase speed;

These posts, in common with other areas outside of experimental HEP, have historically been more difficult to obtain, as by definition the work is always for long term benefit and not immediate operations.

The Vertical Layer can in principle be requested through grants (PPGP or PPRP) where it is experiment specific.

For any WLCG wide work it can in principle be requested through GridPP.