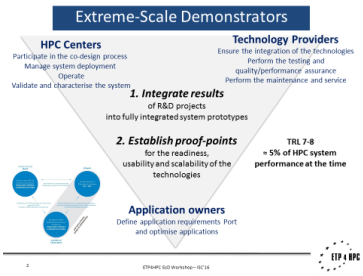

ETP 4 HPC
THE EUROPEAN TECHNOLOGY PLATFORM FOR HIGH PERFORMANCE COMPUTING

The Extreme-Scale Demonstrators concept: Current State of Definition

Follow this discussion at: http://www.etp4hpc.eu/en/esds.html

1

1 Preface

The scope of this document is to serve as background information for HPC stakeholders and HPC user communities interested in learning more about ETP4HPC's concept of "Extreme Scale Demonstrators" (EsDs. The objective is to formulate a clear and compelling proposal to be integrated into the HPC 2018 Work Programme of the European Commission. This document summarises the state of ETP4HPC's Extreme-Scale Demonstrators concept following dedicated workshops to identify the key objectives and challenges:

- The EsD work session during the European HCP Summit week in Prague (9th to 13th May 2016)
- The workshop during ISC'16 in Frankfurt (23rd 2016)

It should be viewed as a 'living document' with an emphasis on capturing all relevant thoughts and suggestions in preparation of ETP4HPC's next Strategic Research Agenda (SRA) update. There might be slight overlaps and redundancies between different chapters which were written by experts leading various subject fields in the EsD working group.

The next step will be at the EXDCI workshop on 21st and 22nd September in Barcelona. Its objective is to:

1. Extend the use case discussion and analysis (see Chapter 4) onto additional science communities both in the High-Performance Computing (HPC) and in the HPDA (High-Performance Data Analytics) arena.
2. Facilitate a discussion with system integrators on their views and suggestions in mobilising critical resources for a successful EsD programme implementation.

2

2 Extreme-Scale Demonstrators (overview included in SRA 2)

High-Performance Computing (HPC) is a crucial asset for driving European innovation and strengthening its technology providers. The "Extreme-Scale Demonstrators" (EsDs) proposed are vehicles to optimise and synergise the effectiveness of the entire HPC H2020 Programme through the integration of isolated R&D outcomes into fully integrated HPC system prototypes; a key step towards establishing European Exascale capabilities and solutions. The primary focus of the EsD projects will be in establishing proof-points for the readiness, usability and scalability potential of the successful technologies developed in WP2014/15 and WP2016/17, when deployed in conjunction with open market technologies at that time. Accordingly, this presents ETP4HPC's current view in the way these EsDs should progress.

There is an existing consensus between HPC centres and industrial members of the ETP4HPC that such projects should create "ready to use" systems commensurate with Exascale commercial objectives. They should encourage a strong co-design approach between technology and applications providers. They would then produce tangible results which validate the capabilities produced in the preceding H2020 work programmes. These EsDs should provide platforms deployed by HPC centres and used by CoEs for their production of new and relevant applications. The fully integrated EsD systems should not be confused with systems/subsystems prototyped as part of individual research projects but as synergetic

and integrated platforms. At the project end, the EsDs will have a higher TRL[1] (Technical Readiness Level) of 7-8 (compared to 6-7 of prototypes as part of the projects), thus, their stability and usability will enable stable application production at a reasonable scale[2]. Therefore, the EsDs will be 'stepping stones' towards a more expedited and solid commercial exploitation of the underlying system design and technology. The subsequent commercial exploitation of EsD output in preparing exascale level products and/or component technologies would be left to the participating industrial partners, whilst in relation to the integration of EsDs, the target is to deploy technology developed in the FETHPC programmes. The EsD projects need to be open to also include relevant technology developed outside of this programme, i.e. in the other parts of the H2020 programmes or in the global market.

It should be emphasised that the purpose of creating and using the EsDs is fundamentally different from procuring 'large commercially available production systems' e.g. by Tier 0 centres within PRACE. The EsDs are meant to validate and prove the advancements in R&D performed within the H2020 HPC work programmes and to gather valuable feedback for future projects, whilst the periodically procured commercial HPC systems are geared towards providing a robust compute infrastructure available to large user communities. However, a fraction of cycles from the EsDs, once deployed and stable, should be made available to members of larger user groups, with well-established allocation mechanisms to also expose the technology to the wider community. The EsD projects should, therefore, help industrial (and also SME) users to prepare for the next step in their HPC usage.

In summary, the EsD projects will fill the following important gaps in the current HPC H2020 programme:

- Bringing the technologies developed in the FET-HPC programme and related H2020 R&I closer to commercialisation, thus, fostering exploitation and take-up of these technologies; 3
- Benefiting from targeted R&D efforts across many projects and to combine components into an integrated system;
- Providing the missing link between the three pillars of the HPC strategy: technology providers, infrastructure providers, and user communities through projects that leverage their respective expertise to develop new high-end compute platforms.

2.1 Approach

ETP4HPC proposes that EsD projects will be set-up as 'dedicated R&I projects' within WP2018/19, maximising the level of co-design evolving since WP14/15. The current view of ETP4HPC is to have two sets of EsD calls, each one leading to one or two projects. Each of the EsD projects will be structured in two phases:

[1] Metrics used according to http://www.acq.osd.mil/chieftechnologist/publications/docs/TRA2011.pdf

[2] A performance target and size around 5 % of the peak systems at that time is recommended.

- Phase (A): Development, Integration and Testing, involving little or no basic technology research projects, will have a substantial R&D focus mostly geared towards integrating and customising hardware and software components and sub-systems developed in the preceding R&D projects.
- Phase (B): Deployment and Use, where the EsD is validated and operated by a hosting centre and made available to application owners for code porting and development to address numerical/extreme data challenges as well as characterisation and platform validation based on real use cases.

The following Figure 1 shows the relationship between the EsD projects and the work programmes WP2014/15 and WP2016/17:

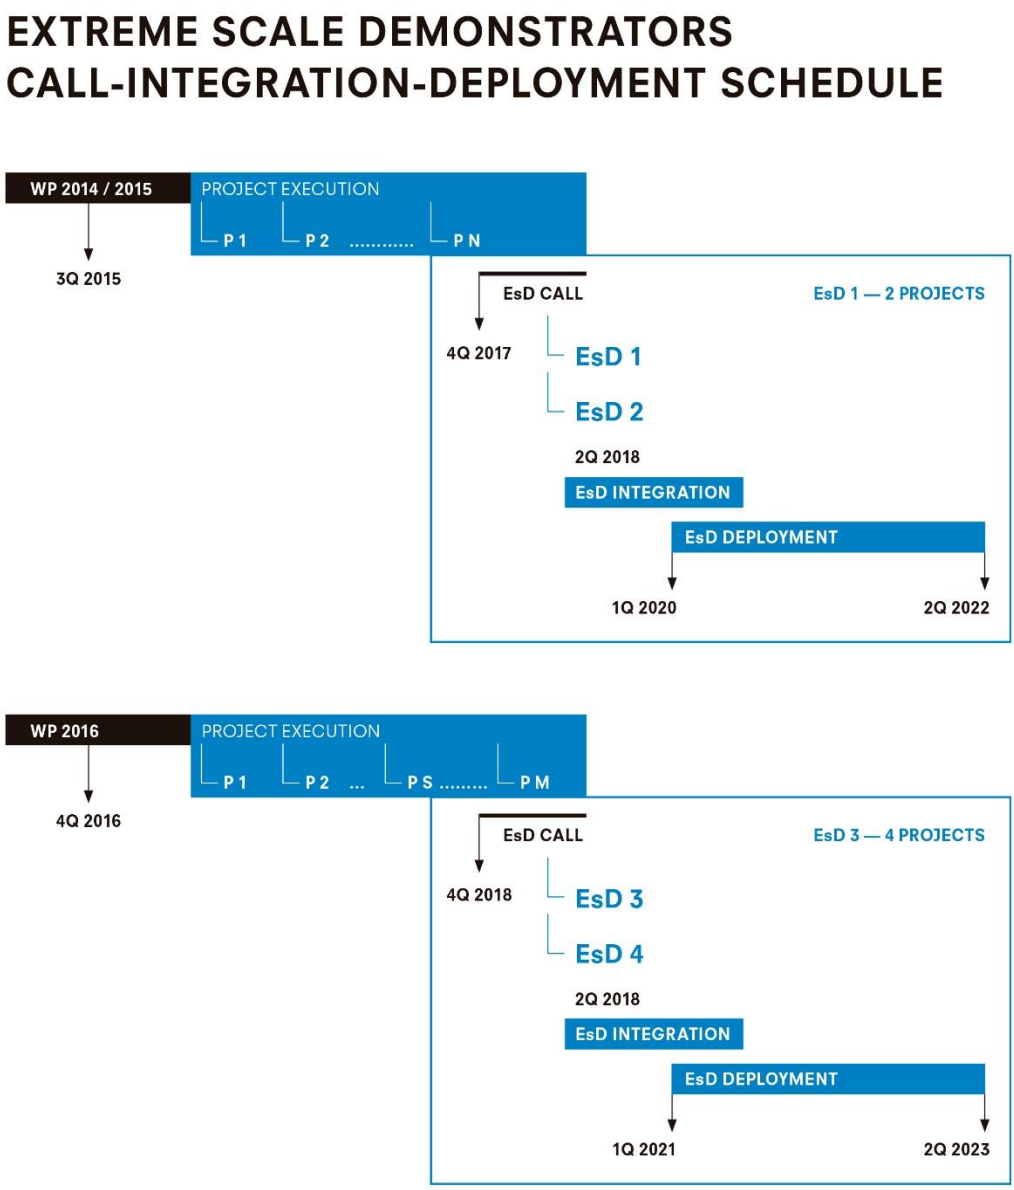

Figure 1 - The relationship between the EsD projects and Work Programmes WP2014/15 and WP2016/17.

ETP4HPC suggests scheduling a block of EsD projects immediately after each of the first two work programmes. Each block should fund between one and two EsD projects, according to the available budget and the co-funding structure.

2.2 Proposal of ETP4HPC for the EsD calls

The EsD calls will have a high dependency on the outcome of WP2014/15 and WP2016/17 projects regarding their timing, but mostly regarding contents. The portfolio of accepted projects in the work programmes must provide a sound technology basis for building EsDs, and the accepted projects should be actively encouraged to foster cross-project interlock. The structure of the WP16/17 should support cross-project integration, with particular regard to IP visibility and licensing clarity. Little coherence between accepted projects, too many disjoint focus areas, and insufficient technology options and readiness might otherwise jeopardise the success of the EsD calls.

It is proposed that calls should be announced within the WP 2018-2019. It is proposed that the EsD project calls will have a funding envelope compatible with a spending of €20-40M (30-50% R&D and 50-70% parts costs) per EsD project for phase A and €3-6M for phase B to cover utilities, operation-manpower and maintenance. Phase A should have a duration of 18-24 months and phase B of 24 months with a validation feedback checkpoint after 9 months. Therefore, total project duration of 32-48 months is envisaged.

The EsD characteristics will need to be further refined, however, they should deliver a high enough TRL to support a stable and effective production environment in their respective Phase B. Their impact on commercial product lines is not expected before 2020. Looking at the hardware characteristics, it is expected that the EsD architectures target scalability of applications up to 200 Pflops. This and other hardware characteristics (energy efficiency, I/O bandwidth, resiliency, etc.) will be detailed in the 2017 release of the SRA, also taking into account results from the FETHPC projects and requirements from the CoEs.

5

ETP4HPC recommends suitable projects to involve three types of partners for EsD projects: technology providers, application owners and HPC centres.

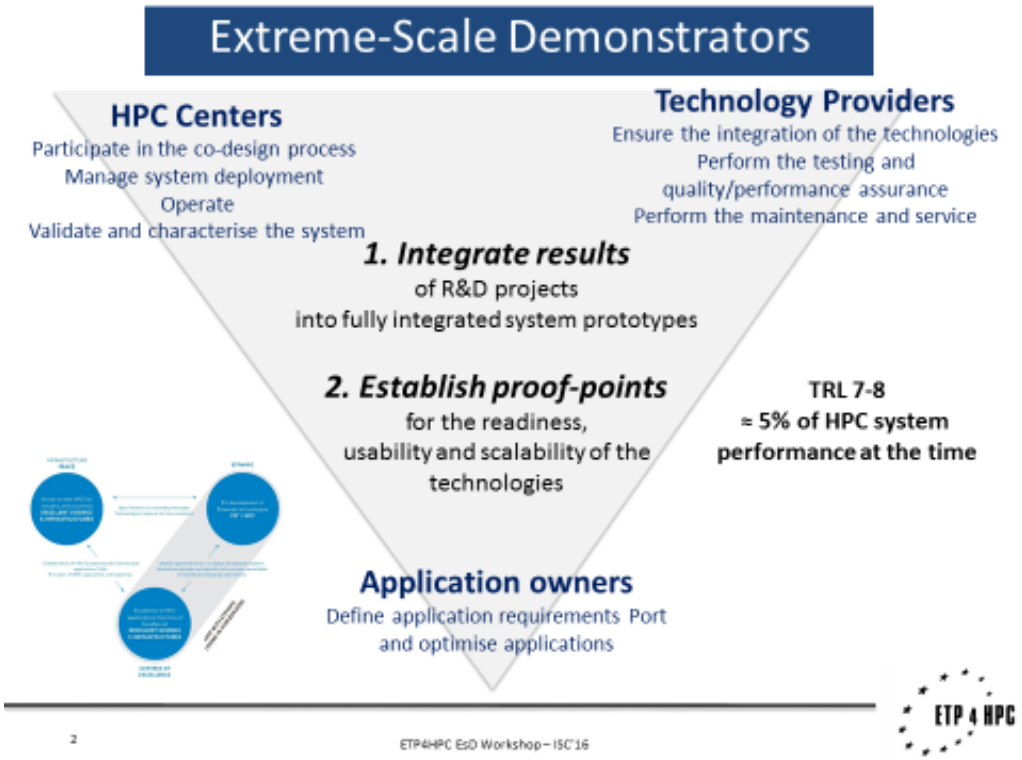

Figure 2 - The roles of organisations representing the three pillars of the European HPC Eco-system in the development of the EsDs.

6

The role of **technology providers** (with a key role for system integration) will be to ensure the integration of technology, the project management, the testing and quality/performance assurance during phase A. The system integration will be the focal point for maintenance and service during phase B.

The role of the **application owner** will be to define application requirements and key challenges, which requires Tier-0 type resources, and that can be addressed by EsD during phase A. During phase B, they will port and optimise application(s) to EsD and use EsD productively.

The role of participating **HPC centres** will be to participate in the co-design process and to manage system deployment during phase A. Furthermore, they will operate the EsD, validate and characterise the system prototypes (in terms of performances, robustness, efficiency, etc.) during phase B.

3 Meta-Level Workload Requirements & System Characteristics (Hans-Christian Hoppe)

One of the key objectives of the call for EsDs is to foster the creation of system prototypes which serve the needs of the relevant usage areas in the European Research Arena and for European industry, both for traditional HPC simulations and emerging data-intensive applications. Understanding the usage scenarios and their commensurate requirements for Exascale-class systems will enable a data-driven discussion on the advisable number and orientation of EsDs to best cover these requirements.

This does constitute "the upper half" of a co-design cycle, as depicted in Figure 3. The main objective is to identify the relevant usage scenarios, characterise the underlying application(s) and derive their specific requirements for future system properties at a high level of abstraction. It is important here not to prescribe the implementation of systems in detail – it will be the task of the proposers of EsDs to decide which part of the requirements space to address, and how to architect and design the respective EsD.

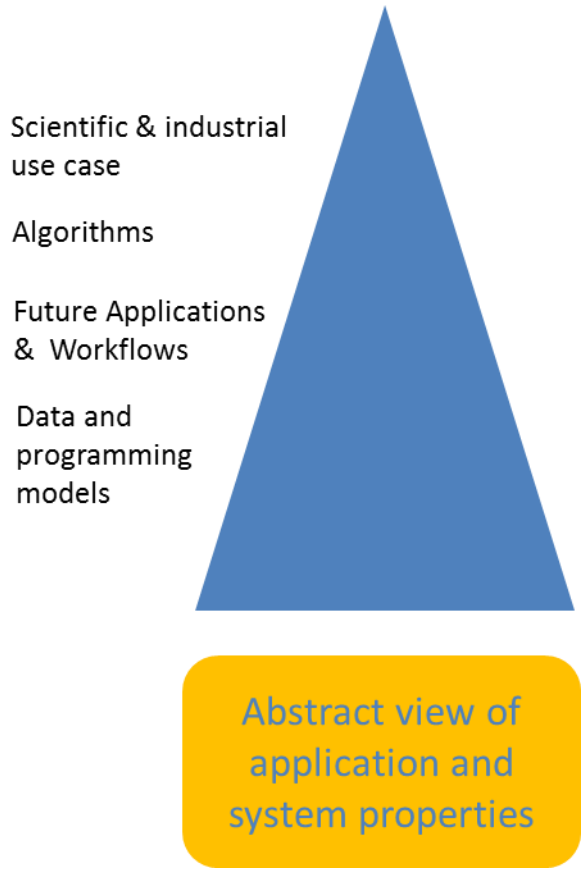

Figure 3: Deriving the application requirements and system properties for EsDs

The components of the abstraction layer can be summarised as:

- Use case/application information – such as scientific or industrial domain, algorithms used, scalability properties and objectives, simulation vs. extreme data case etc.
  This information is critical to assess whether we have indeed covered the space of relevant use cases.
- Quantifiable, specific system properties required by the application(s) – such as integer and floating point performance, interconnect performance, memory, interconnect and I/O access patterns etc.
  This information is required to identify Clusters of system properties and assess the fit of EsD proposals.
- Non-quantifiable system properties – like programming models etc.
  This information is useful to guide EsD proposers on how to layout their systems.

The next sections discuss the above categories in more detail. They are based in part on the work done by Geoffrey Fox (Indiana University) on application classification to assess actors in HPC/Big Data convergence (http://web.cse.ohio-state.edu/~luxi/hpbdc2016).

3.1 Use Case/Application Information

To provide high-level insight into the properties of the application(s) that define the use case, information on the data structures and the algorithms operating on these are required:

- Data structures: n-D structured or unstructured meshes, trees, graphs, particles, statics vs. dynamic (e.g. adaptive mesh refinement), …
- Computational algorithms used: dense or sparse LA, particles & fields, n-body methods, spectral methods, multiscale methods, explicit vs. implicit solvers
- Data algorithms used: streaming vs. "in place", search/query/indexing, ML methods used/planned, ratio of operations vs. transactions

For the EsD and the Exascale timescale, the scaling method and targets of the use case and the parallelism inherent in the application(s), as well as the inherent or intended parallelism model should be given:

- Scaling of the use case: strong vs. weak vs. ensemble/farming
- Parallelism: realistic degree of parallelism (# of threads and processes, # threads per process) inherent in the application(s) and ways it is expressed in, such as embarrassingly parallel, bulk synchronous, dataflow/workflow, map-reduce (general)

3.2 Specific, Quantifiable System Properties

From a more detailed characterisation of the use cases and applications, specific and quantifiable system properties can be derived. These will be key in identifying Clusters in the system design space, and in matching EsD proposal with use cases. As above, it is important to think ahead into the EsD and Exascale system timeframe and to anticipate the evolution of use cases and applications.

8

- Delivered computational performance for the whole system and for processes/threads, in terms of floating point (DP, SP, lower precision) and integer operations. Focus should be on the full system performance.
- Required memory sizes for full system; scaling of memory sizes per thread or processes depending on the number of same.
- Memory performance (latency, sustained bandwidth), compute vs. memory ratio (Byte/Flop ratio)
- Interconnect performance (point-to-point, bisectional bandwidth)
- I/O volumes and patterns (aggregated and evolution over runtime) as well as required I/O bandwidth
- Energy to solution by the full system (KWh), maximum power use per compute unit (KW), energy per operations (pJ/Flop, pJ/communication etc.)

3.3 General Non-Quantifiable System Properties

Finally, a number of important characteristics and requirements of applications cannot easily be expressed as quantifiable metrics. Yet, they are important in putting the quantifiable data above into perspective, and providing valuable information to guide the system architecture and design of theEsDs. These do include:

- Programming models: OpenMP, task-based dynamic threading, MPI, PGAS, distributed data flow, etc.
- Execution models: batch single-step, workflow multi-step, interactive viz/control, malleable execution, etc.
- Communication structures: n-D halo exchanges, any-to-any communication, reductions, scatter-gather
- I/O structures: relation of processes vs. files, prevalent read/write/append mode, sequential vs. random access, etc.
- Memory access patterns (spatial & temporal): contiguous/strided/indexed, data reuse, repeating address streams
- Heterogeneity: steps/phases in the application with different requirements (please provide separate category #1 and #2 data for each)

4 Use case analysis: ISC workshop outcome (Erwin Laure)

In a first EsD workshop in Prague during the HPC summit the need for a user-driven approach for EsDs was clearly stressed. The following high level recommendations were articulated:

- EsDs need to be driven by scientific challenges and the application characteristics should drive the system design.
- Co-design from the very beginning is important. A model, where applications only come into the game at a late stage is not considered viable. Both academic and industrial challenges should be considered.
- The success metrics of an EsD should not be Linpack flops but the science it enables (*to be further elaborated*).
- EsD proposals should be required to identify the scientific challenges they will contribute to and how this contribution is structured.
- Application porting and optimization needs to be appropriately funded. This could either happen as part of the EsD project (with appropriately increased overall funding), or through other means outside the project. The latter bears the risk that system and application work will not be in sync.
- Funding for applications needs to be limited to porting and optimization**.** Method development or development of new applications should be funded elsewhere. Science should not be funded either, but will be supported through free cycles.
- EsDs need to consider full applications (and thus need to have the appropriate system ware available) as well as kernels/mini-aps. The former is needed to demonstrate how the EsD will contribute to solve scientific challenges; the latter is needed for testing and development purposes. Hence it is likely that the EsD will have a "production" and "development" "partition" (in whatever form that it will be realised). They need to provide an appropriate QoS level.

- EsDs should have a longer lifetime, probably e.g a standard production system to make the investment of them viable.

Following these initial discussions, a number of user communities started to discuss their approach towards EsDs at a workshop at ISC, 2016. Following an open call for expressions of interest, 6 communities provided input to the workshop:

- Astronomy (SKA)
- Biomolecular Simulations (BioExcel CoE)
- Climate (ESiWACE CoE)
- Energy (EoCoE CoE)
- Materials (MaX CoE)
- Bayesian Uncertainty Quantification (ETHZ)

It is important to stress that this set of communities is by no means a comprehensive set of communities which EsDs might target, but merely a starting point composed of communities that were able to provide input at the time of the workshop. Still, initial input could be collected from a relatively diverse set of communities.

The scientific case for Exascale computing has been provided by previous studies, including the PRACE scientific case[3] or the ESSI project reports[4]. In this workshop we focused on gathering initial application input on potential EsD system characteristics. This exercise can be seen as a starting point for future discussions. Concrete EsD proposals will need to have much deeper discussions between applications and system design. Co-design will be an important tool for this as both, applications and system design need to evolve together and both might change depending on the characteristics of the other.

We have summarised the findings of the workshop here below, the full input of the communities can be found at http://www.etp4hpc.eu/en/esds.html.

- ***Floating Point/Node performance***
  The need for powerful nodes was stressed by many domains. This includes large vector unites but does not necessarily mean double precision as several domains could also work with lower precision. Trading node performance with scalability was not considered to be a viable way forward. Efficient management of heterogeneous components (e.g. CPU/GPU) will be required.
- ***Network latency***
  Low network latency is important for many applications. Bandwidth, given the current state of affairs, is less of a concern, however, still important. Of course, data intensive applications still need high bandwidth. The network should also provide performance counters etc. to allow a better application-level understanding of its performance. Load balancing and asynchronous communication will be important.
- ***Memory bandwidth/performance***
  High memory bandwidth is important particularly for data intensive applications. Systems should have a high flops/bandwidth ratio.

[3] http://www.prace-ri.eu/prace-the-scientific-case-for-hpc/
[4] http://www.eesi-project.eu/wp-content/uploads/2015/05/EESI2_D3.3_Final-report-on-application-grand-challenges.pdf

- ***I/O***
  I/O is considered important for most applications, particularly for check pointing and enabling workflows.
- ***HTC/Workflows***
  Many applications employ ensemble computing or complex workflows combining HTC and HPC. This poses issues on I/O as discussed above and also on the system management software (particularly batch schedulers).

Other Considerations
The system will likely have a power cap for financial reasons (e.g. SKA expects to be able to support a system using up to 5MW).

The system will need a versatile software stack including all commonly used programming environments (MPI, OpenMP, CUDA, etc.) but also virtualisation, scripting, and performance/correctness tools. Applications will need to scale well beyond 1M cores.

Measures of success for EsDs will need to be defined – pure Linpack performance is not meaningful for real applications.

5 EsD project budgets: Maximising the Value of ESDs (Malcolm Muggeridge)

5.1 Budget Goals

The proposal aims to maximise the value for Europe of all the ESDs created in the program. To do this it is believed that there needs to be the ability to make an informed selection of the characteristic capabilities of all proposed ESD proposals ability to meet the goals, beyond a simple comparative selection through evaluation.

This is driven by the belief that one solution is unlikely to meet all the needs of future Extreme scale systems and that to maximise future European competitiveness it is essential that there is both strength and depth to the European systems capabilities.

If this is not done there is likely to be high levels of overlap in the capabilities of proposals, a focus on potentially narrow use cases, and lack of inclusion of other equally important capabilities necessary to ensure wide coverage of the wide ranging needs for Europe.

5.2 Evaluation Process

A “two-staged evaluation process” is proposed where initial proposals are presented but within them a level of flexibility is requested such that a core system can extend its capability in either ability to address more use cases or to increase the performance of specific use cases.

In the first stage of the evaluation the goal will be to have maximum coverage of a defined set of high level system characteristics (see also chapter 4) across all the proposals provided and within those proposals allowing the evaluators to optimally select not just the overall proposal but provide feedback

to the proposers of goals that may be included/expanded or excluded/retracted. A best cost/coverage method can be used to optimise within the overall budget.

Conceptually these system characteristics can fill a 'Spider diagram' of multiple dimensions of ESD goals with a complete circle of capabilities at an optimised cost to the program.

In the second stage the 'refined' proposals will be evaluated and the full cost of each can be optimised.

5.3 Proposal Structures

It is proposed that each project will present in stage 1:

a) A Core systems describing
   a. Its capability to meet one or more of the stated call goals
   b. The associated use cases/applications that will be used to demonstrate the goal(s) can be achieved.
   c. The Hardware and Software elements of the proposal
   d. Any '3rd party'elements of the solution required
   e. The budget required for the system

b) Additional 'Incremental' goal coverage (on an individual or group basis)
   a. Incremental capability to meet further goals
   b. Additional use cases employed
   c. Additional or expanded hardware
   d. New Software elements
   e. Additional 3rd Party elements
   f. Incremental budget required

12

Notes:

- Any number of 'Incremental' goal coverage may be included in a proposal
- If the EsD programme is implemented using two calls (2018 and 2019), proposers in the 2018 call are not restricted to make 'Incremental' goal proposals in 2019.

5.4 Call Budget

The budget breakdown is designed to enable selection of a multiple of alternative core proposals to be adopted but allow the stage 1 selection to maximise the 'incremental' elements across these proposals.

The formula below would allow up to 6 core proposals to succeed but this will reduce as incremental goals are included.

- The overall budget proposal is for 200M€
- Of this it is proposed that 60% be in 2018 and 40% in 2019
- Core proposals should not exceed 25% of the annual call budget
- The total cost of the 'Core' proposal and any number of 'Incremental' goal capabilities should nominally not exceeding 45% and maximally 65% of the available annual call budget.

- A minimum of 2, max of 4 core projects accepted in 2018
- A maximum of 2 core projects in 2019

5.5 Cost Inclusion

All cost elements of the systems deemed necessary should be included:

Costs will be in 3 categories

- Fixed (NRE/development costs).
- Variable with capability (e.g. Performance/scale/capacity of system) – including Hardware Bill-of-Material (BoM) and integration costs, the system scale proposed will be the decision of the proposers and must be justified based upon the use cases proposed to demonstrate the goals.
- Variable with operation, i.e. power, system administration and maintenance.

Some other elements to be included:

- Cost of any software purchases or licenses required to operate the system.
- Infrastructure costs of physical location, any hosting or connectivity.

6 EsDs: Options for acquisition (Dirk Pleiter)

13

This chapter presents different options to realise the acquisition[5] of all equipment required for an Extreme-scale Demonstrator (EsD). This document is intended to be a working document collecting current knowledge about different options, their strengths and related opportunities as well as their weaknesses and related threads. Note that we consider system acquisition only. How to handle costs related to the deployment[6] and operation of an EsD will be analysed elsewhere.

We will first document relevant boundary conditions before analysing different acquisition options. The list of options is not considered to be complete as, e.g., also models involving renting or leasing may be considered. Finally, we look into the options for a light-weight and flexible funding scheme aimed at promoting involvement of SMEs.

6.1 Relevant boundary conditions

The current EsD concept involves 3 different types of actors:

- **Technology providers** have the role of providing hardware and software technologies that become part of the EsD. One of them has the (possibly additional) role of **system integrator**, who is

5 In this document we use the term “acquisition” in order to be able to give more specific meaning to the terms “procurement” and “purchase”.

6 It can be expected that at least part of the deployment costs could be included in the acquisition costs.

responsible for building the EsD and later deploying it at the hosting site. In the following we will assume the system integrator as being a commercial operator.

- **Infrastructure providers** are the experts for operating large-scale compute infrastructures. While there might be multiple infrastructure providers involved in an EsD project, only one of them will act as **hosting institution** for the EsD. In the following we will assume the hosting institution to be a public entity and thus a contracting authority according to Directive 2004/18/EC.
- **Application providers** are experts for science and engineering applications that will productively exploit the EsD. They are assumed not to play any role in the process of acquisition of the equipment.

It is currently proposed to aim for 4 EsD projects with the first projects not starting before end of 2018. The targeted time for deployment of the first EsD systems is, therefore, in 2020.

The realisation of an EsD will generate a significant amount of costs related to acquisition, deployment and operation of such a system. Furthermore, the current EsD concept assumes that the application providers are not supposed to be attracted by personnel resources provided through an EsD project, but rather to be attracted by a vast amount of computational resources that will enable them to address science and engineering challenges, which otherwise would not be possible at this point of time. As a consequence it will be important to keep the EsD operational for a longer period of time, i.e. for at least 3 years and preferably for 5 to 6 years.

6.2 Option A: Combination of PCP and PPI

Design and acquisition of an EsD could be realised by combining a Pre-Commercial Procurement (PCP) and a Public Procurement of Innovative solutions (PPI). A PCP is a procurement scheme that allows public institutions to procure R&D services, which can additionally be financially supported by the European Commission (EC). It is organised as a multi-phase procedure where the last phase involves the development of first products. These products could then be procured through a PPI, which is a co-funding model of the EC to stimulate procurement of innovative solutions at an early point of time. The procurer thus acts as lead customer.

PCP and PPI are based on well-defined processes. However, they are not considered to be suitable for acquisition of an EsD. The processes are not compatible with the envisaged schedule for deploying EsDs in 2020 as the execution of a PCP will require at least 3 years. Both, PCP and PPI can within H2020 only be co-funded by the EC and thus a significant fraction of the funding would need to be secured from other sources. The process would require some actors to become procurers while others become competing suppliers. This is not compatible with the envisaged mode of collaboration. Finally, the number of suppliers needed for the initial phase of multiple PCPs might have to be too large compared to the potential number of bidders.

6.3 Acquisition of goods within an RIA project

In case of EsD projects being executed as Research & Innovation Action (RIA) the costs for an EsD might be claimed as equipment costs as the beneficiary may purchase goods necessary to implement the action (see [H2020-MGA], article 10).

The H2020 Model Grant Agreement (MGA) does, however, foresee the following conditions:

- Clause 10.1.1 [H2020-MGA] states: “The beneficiaries must make such purchases ensuring the best value for money or, if appropriate, the lowest price.” In the comments it is noted, that “best value for

money principle does not require competitive selection procedures in all cases", but it must be demonstrated how best value for money was ensured.

- Clause 10.1.2 [H2020-MGA] states: "Beneficiaries that are 'contracting authorities' ... must comply with the applicable national law on public procurement".

While H2020 foresees a funding rate of 100%, there are possibly limitations for claiming all costs that occur in the context of the acquisition of an EsD. By default, only depreciation costs of equipment, infrastructure or other assets can be claimed (see [H2020-MGA], clause 6.2.D). The duration of the project must thus be carefully chosen to ensure that all costs can be covered. It is worthwhile noting the the MGA does foresee exceptions to this rule if it is foreseen in the work programme.[7]

6.4 Option B: Public procurement by hosting institution within RIA

As the hosting institution is assumed to be a contracting authority it is bound to public procurement rules, i.e. acquisition needs to be performed through a procedure that obeys the principles of public procurement: the procedure must allow for fair and open competition and be conducted in a transparent manner. As the goal is to acquire technologies that have been created and integrated by project partners, suitable approaches have to be identified for restricting competition while ensuring best value for money being obtained. Restrictive procedures (see [2014/24/EU], article 28) or innovation partnerships (see [2014/24/EU], article 31) may provide the right opportunities.

6.5 Option C: Purchase by system integrator within RIA

This option foresees equipment to be procured by the integrator (and/or other commercial operators within the EsD consortium). The integrator ensembles the system and provides it to the hosting institution on loan.

While it can be assumed that all relevant integrators have internally well-established purchasing processes, there are nevertheless various challenges. For small integrators the financial risks for building a larger system might become too large as they can only claim costs later. This might lead to consortia becoming dominated by larger commercial operators. Another problem, which has been identified, is possible restrictions for suppliers acting as integrator, who wish to purchase equipment that they produce themselves.

7 This might be a topic of discussion with the EC when providing input for the next work programme.

6.6 Comparison of options

The following table summarises the analysis of the different options:

	Option A	Option B	Option C
Strengths	Based on innovation driven process	Well established public procurement procedures are in place	Leverages existing and simple processes
Weaknesses	Duration incompatible with EsD timeline; significant funding from other sources needed	Challenge of keeping procurement open for competition	High financial risks for small integrators and possible dominance of larger suppliers within consortium; integrator cannot purchase from himself
Opportunities	-	Exploit new instrument of innovation partnerships	-
Threads	Too small a number of suppliers that could enter a competitive PCP process	Procedure can be challenged by suppliers outside that are not part of the EsD project	-

6.7 Options for encouraging SME involvement

SMEs might make important contributions to EsD projects in different ways. They could contribute technology that is integrated in the EsD. Furthermore, SMEs could investigate and demonstrate simulation services on EsDs. Their involvement, however, might be hampered by financial risks and the lacking ability to commit to a multi-year project. For EsD projects it could thus be beneficial to foresee mechanisms specifically targeting involvement of SMEs.

A possible template for promoting SME involvement has been established by the EC-funded FORTISSIMO project [FORTISSIMO]. It regularly calls for proposals to perform experiments on HPC systems. The mechanism is light-weight as the length of proposals is limited to 10 pages and proposals can be submitted electronically. Successful applications are funded by up to €250,000 and applicants join the project as third party. See [FORTISSIMO-D72] for more details on the procedure.

6.8 Related aspects

In the following list we collect a set of aspects, which are related to the choice between the different options for acquisition:

- Different options for acquisition result in different ownership relations, which might impact access restrictions to the system. This could, e.g., become relevant when the EsD is exploited by a commercial operator, in particular if it is not a member of the consortium.

- Operation of the system may require that support contracts and license agreements are in place. How this is managed may also depend on how the EsD had been acquired and who formally owns the system.

6.9 References

[2014/24/EU] European Parliament and European Council, "Directive 2014/24/EU of the European Parliament and of the Council," 26 February 2014.

[FORTISSIMO] http://www.fortissimo-project.eu/ (accessed 16 July 2016).

[FORTISSIMO-D72] Guntram Berti, Guy Lonsdale, Francis Wray, "FORTISSIMO D7.2 Open Call 2 Report," 29 September 2014.

[H2020-AGA] European Commission, "H2020 AGA – Annotated Model Grant Agreement," version 2.1.1, 1 July 2016

7 EsD Consortium – a more detailed view (Thomas Eickermann)

Given the objective of the EsDs to bring HPC-related technologies developed in FP7 and H2020 projects closer to commercialisation, using a strong co-design approach, a successful EsD project needs to include competences (capabilities and capacities) from all of the three pillars of the EC's HPC strategy: technology providers, infrastructure providers, and user communities. The following paragraphs summarise the current state of discussion on their respective roles and the implications on project consortia.

7.1 Technology Providers

EsDs will include innovative technology developed in EC-funded R&I actions and possibly elsewhere. The integration of these technologies will require further R&D to be conducted within the EsD projects. Therefore, providers of such technology need to be members of the consortium. It is expected that in order to build a complete system, inclusion of additional state-of-the art components and technology will be required. Such components can be bought from the market, i.e. from contractors external to the consortium.

It should be noted that these considerations apply both to hardware and software technology. Formerly, ongoing and future FET-HPC projects and similar R&I actions provide a substantial pool of suitable technology. Integrating technology developed in unrelated projects thus far will be challenging. Therefore, interested parties are encouraged to enter discussions about forming consortia as early as possible, in any case well before the publication of a call.

Among the technology providers, the **integrator has a particularly important role**: in phase (A) of the project, it will lead the development, integration and testing of the EsD based on technologies and components provided by different parties. In phase (B) it will have a leading role in stabilising, improving and maintaining the system. To take this responsibility will require a company with the capacity and experience to develop and deliver large HPC systems.

It is expected that SMEs will be important partners in consortia. They are often highly innovative and agile, and can contribute e.g. specific components to the demonstrators. However, participation in a high-risk R&I action is particularly challenging for an SME with its limited financial capacity. Therefore, specific actions may be considered to encourage the participation of SMEs. This could e.g. include lowering administrative burdens.

Finally, a **system architect** (an individual or a small team) will be needed to translate user requirements into system specifications during the co-design process.

The approach to have the integrator and the majority of technology provided by consortium members also has an impact on the procurement model and ownership of the EsD. It will typically lead to some form of joint ownerships among the contributing consortium members and the consortium needs to develop careful agreements on the final ownership of the demonstrator as well as on the newly developed IP.

7.2 Hosting Centres

The main role of the hosting centre during the project is to deploy and operate the EsD and make it available to application owners for porting and development in phase (B), and, once stable, also partly for regular usage. This will include providing advanced user support. These tasks are fully aligned with the mission, skills, available in the infrastructure of typical large HPC centres.

In phase (A) the hosting centre will also play an important role in the co-design process, both providing operational requirements and leveraging its links to application communities and the experience in mapping user requirements to system specifications.

18

A critical point for the buy-in of the HPC centres to host an EsD will be the funding model. The operational costs for a system of the envisaged size can be several million euros per year. While deploying and testing beyond-state-of-the-art systems is often part of the mission of large HPC centres, operational costs on this level for a demonstrator can constitute a delimiting factor.

7.3 Application Providers

Application providers will also play a key role in the project right from the beginning. In phase (A) they will be part of the co-design process, by contributing a scientific challenge and the corresponding application or set of applications. In phase (B) they will port the application codes to the EsD and contribute to the assessment of its performance and usability. While it is not within the scope of the EsD projects to develop new applications, it is expected that the new or advanced architectures will require substantial development effort for porting and harnessing their potential.

Due to their critical role in both phases of the project, it is mandatory that application developers are part of the consortium. While getting early access to a new technology and a system with substantial capability and capacity will be a huge incentive, it will nevertheless be required to also co-fund the work of the application owners to ensure that the required skills and resources can be allocated in accordance with the project plan (about 5% of the overall project funding has been considered a typical amount). For application developers that have already been in contact with technologies deployed in the EsD, e.g. through the FET-HPC projects, the entry barrier will most likely be lower.

The consortium member(s) that represents the application community is ideally the application developer or has close links to the developer to ensure that the development work performed is fed into the main application code-base. While CoEs have the required skills, they often have limited free resources. In addition, a CoE member in the EsD consortium could help to create strong links to the CoEs. Furthermore, the EC could consider allocating resources in future CoE calls for cooperation with EsDs, and also, an SME providing support for an OpenSource code could potentially represent an application community in the consortium.

7.4 Project Coordination

While the selection of the Coordinator is a choice of the consortium, some considerations about suitable options have been discussed. The integrator and the hosting centre will have the key role in phase A and B of the project, respectively. This makes them primary candidates to also take on the role of project coordinator. The integrator as Coordinator would emphasise the ambition to raise the TRL of the EsD to a level that allows productisation. On the other hand, the natural links that a hosting centre has to all partners (technology providers and users) are beneficial for a Coordinator. Furthermore, the hosting centre has no commercial interest which might lead to potential conflicts of interest.

A third option is that the integrator and hosting centre take the technical lead in phase A and B respectively, and a third consortium member takes the role of the Coordinator, focussing mainly on the contractual and administrative tasks.