

ETP4HPC ESD Workshop, Prague, May 12, 2016

Interactive Session Input

Technical Characteristics of the ESDs.....	2
Required ESD project budget assumptions.....	3
Procurement model options	5
Composition of consortia-roles and tasks.....	6
Use cases Enabled by EsDs.....	8
Additional Material	9

Summary:

Technical total system target characteristics

- Peter Hopton
- Igor Zacharov

Required ESD project budget assumptions

- Malcolm Muggeridge
- Jean-Pierre Panziera

Procurement model options

- Dominik Ulmer
- Piero Altoe

Composition of consortia-roles and tasks

- Hugo Falter
- Estela Suarez

Use cases Enabled by EsDs

- Mark Bull
- Alan O'Cais

Technical Characteristics of the ESDs

- Peter Hopton
- Igor Zacharov

1st Round

Reporter: Peter Hopton

1) What is Exascale

- a. ExaFlop or Exabyte or ExaAnything
- b. Not ExaWatts definitely less than 60MW

2) Size to demonstrate scalability. All about Scalability of:

- a. Compute
- b. Storage
- c. Interconnect
- d. Software
- e. Quality
- f. Reliability
- g. Resilience
- h. Infrastructure
- i. Fault Tolerance
- j. Management/maintenance

3) So it should be Holistic, or End to End, i.e.:

- a. Software
- b. Tuning
- c. OS
- d. Hardware
- e. Management
- f. Infrastructure
- g. Facility

4) Should be equivalent in capability to current top tier systems, but demonstrating (total) cost or energy benefits for scaling

5) It should allow for flexibility but be able to show focus, eg.

- a. Qos
- b. Memory Bandwidth
- c. Compute
- d. Data Centricity/IO
- e. Application

6) The ESD should be productive

7) It may wish to consider openness of interaction for enhancing its capability over time

8) May wish to cap energy at 1MW, including the facility

9) May wish to consider the reuse of waste heat

2nd round

Reporter: Igor Zacharov

Metric: Real Performance

- Comparison with best (commercial) machine at that time

What is “Performance”:

- The **result** (defined by the application) per spend energy for the whole machine (incl. infrastructure) = Power Efficiency
- The run has to be useful (i.e. achieve the scientific goals)
- The machine size is ~1MW electricity consumption
 - Any hardware available at installation time with optimization related to apps
- Application forces machine design and choice of (hardware) optimization
- machine should also be balanced for “product-ization” after EsD
- Reliability is part of the design goals driven by maintenance require.
 - Programmer productivity is part of the community appeal to attract applications to the EsD

Required ESD project budget assumptions

- Malcolm Muggeridge
- Jean-Pierre Panziera

(These minutes combine the report out from both rounds of discussions)

Reporters: Jean-Pierre Panziera and Malcolm Muggeridge

Software Activities in the ESD Context

Intense discussions about the role of software activities (OS, runtime system, tools/programming models, applications) in the ESD context:

- Some parts of the European HPC SW R&D community feel “locked out” due to a perceived focus on HW, and reservations to base their R&D on potentially unstable systems.
- ESDs will of course include at least the system SW aspects, and it is clear that the success of the ESDs depends on SW results.
- Tentative agreement was to that a SW part of 10% in the ESD projects (5 Million for a 50 Million ESD).

ESD Value and Applications

Have to clearly define the value that ESDs bring for the European (science) community. It will not be a xyz PFlop/s performance label alone – scientific results enabled by such a system

should play a major role here. Hence, it will be important to get a good picture of which scientific (and industrial?) applications are the most relevant, and to guide ESD proposals to strive to best support these.

Cost/Budget Breakdown and Scaling

Qualitative cost breakdown for the ESD projects:

ID	Cost Category	Application	Pays for
A	Non-recurring Engineering (NRE)	Hardware, software	People
B	System	BOM and integration	Materials
C	Operation	Power, facilities, operations & support people	Running costs plus people

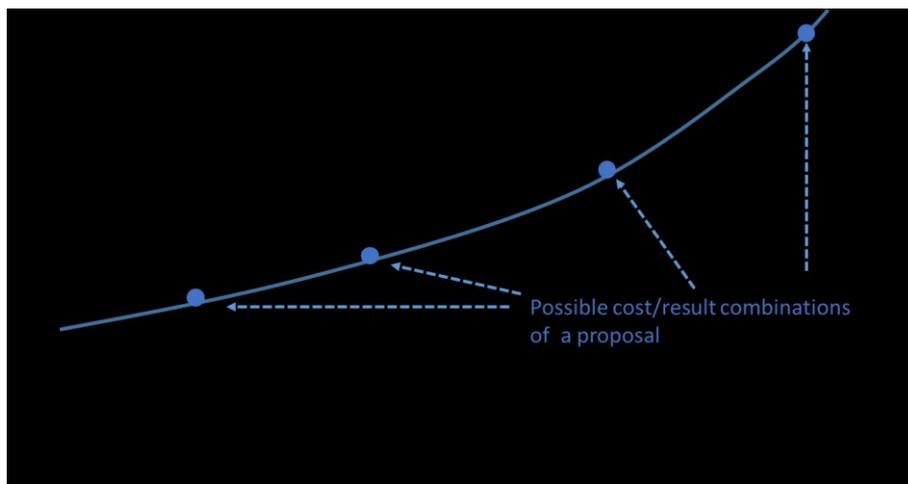
Without precise requirements, no absolute numbers can be given for the categories **A - C**.

The cost categories are expected to show this scaling behavior:

- ESD system size drives categories **B** and **C**.
- ESD technology drives categories **A**, **B** and **C**.
- SW and application work drive categories **A** (system SW) and **C** (applications).

Multiple Cost/Result Points in ESD Proposals

Rather than sizing each ESD proposal to a fixed budget (that in all experience will be identical to the per-proposal budget limit expressed in the call), each proposal could define a number of cost vs. results/value points and let the EU commission evaluation process select which one seems most appropriate. These points can, f.i. differ in system size and capabilities, application domains addressed, uptake of different technology innovations etc. It would be expected that the call would give guidance on the appropriate scales for the results/value.



Such a scheme would give the EU commission the opportunity to reach a “global optimum” in the ESD projects allowed to proceed, i.e. receive the optimum value for these projects and achieve good coverage of all application domains. It also avoids the need to fix number and budget of the ESD proposals at a time when sufficient information is not available.

Number of ESD Projects

While variety and coverage of the relevant scientific domains is absolutely important, it does not seem to be right to fix a number for the ESDs. Taking such a decision would at the minimum require detailed insight into the scope and clustering of requirements across scientific domains, and of the technology innovation options available for the ESD timescale.

The suggestions above would defer a firm decision on number and per-ESD budget to the proposal evaluation time.

Timeline and Project Duration

The alternative to having two separate ESD calls spaced to match the timescales of the H020 FET-HPC projects would be to have only one call for ESD projects, with the option to have projects taking more than 24 months in order to be able to intercept results from the “later” H2020 FET-HPC projects.

Procurement model options

- Dominik Ulmer
- Piero Altoe

Round 1

Reporter: Dominik Ulmer

Questions

Is EsD project resulting in a product?

Who is doing the procurement?

Consortium with one partner doing procurement

Centre + provider team-up

Is commercial partner allowed to make profit?

Can industry use system procured by public entity?

Aspects

Differentiation needed between public procurer and private entity

SME may not be able to act as procurer for larger system

Scenario 1/2

Consortium comprising integrator and HPC center

HPC centre

Bound to public procurement rules

Local procurement rules may allow to procure R&D equipment without tender

Integrator = private entity

Could be procurer like in DEEP project (Eurotech) and Mont Blanc (Bull)

Limitation: integrator cannot procure from oneself

Need to purchase, e.g., from manufacturer

Scenario 1: HPC centre performs procurement

Scenario 2: Integrator performs purchase

Scenario 3

Consortium drafts procurement and looks for external partner

Considered to be difficult

Tendering procedure

PCP is not an option as it is primary for procuring R&D services

PPI is not an option as it is a co-funding model (35% from EC)

Innovation partnerships likely not an option because competition is required

Round 2

Reporter: Piero Altoe

Quick procurement procedures are introduced in UK

Minimum 28 weeks for starting to write tender until awarding contract

Procurement models

Open procurement is not an option

This takes much too long

Does not match to considered framework of consortia

Who is procurer = owner?

EC

Who would evaluate bids?

HPC centre = public entity

Integrator = private entity

Questions

How to manage IPR?

Could procurement be completely be avoid?

This is an R&I action

How can one ensure that procurer=owner has a neutral role?

Distribution of resources

Procurement/purchase of components, e.g. from SMEs

Data centre is possibly in a better position to force opening of system

In particular, if large, industrial company is involved

How to manage heterogeneous concepts for EsDs, comprising, e.g. different modules?

Can Fortissimo can be a model?

All money goes to coordinator, all others are sub-contractors

Fortissimo includes model for proposals

Light-weight procedure allowing for quick decisions

Model could favour SMEs (most responses in Fortissimo are from SMEs)

Other aspects

No funding should be foreseen for data center infrastructure

Private entity must add margin (even if it is small)

If integrator is owner then this does not apply

Need for support contracts

One of the partner has to be formally responsible for that

Composition of consortia-roles and tasks

- Hugo Falter
- Estela Suarez

Round 2

Reporter: Estela Suarez

Coordinator:

- - Two candidates: vendor or hosting HPC centre
- .Vendor would stress the productisation aspect.
- .HPC centre represents users (application communities) and can keep the consortium together: has links to users, vendors, SW providers, and experience on operation.
- - Most participants in the discussion see rather the HPC centre as coordinator.
- - Two open questions:
- .Procurement process issues to be clarified (centre procuring from partner...)
- .Who is covering the operational costs?

Applications:

- - Involve application developers (not application users).
- - Participate from beginning on design process (high level, requirements collection)
- - Co-design already going on current FETHPC projects.
- - No application development effort within EsD, just port and testing.
- - Funding for application preparation to EsD may come from future CoE and FETHPC projects (reserve some part of their funding for it).
- - Also applications are being fund already through current FETHPC and CoE.
- - Expect that applications involved in current and upcoming FETHPC and CoE projects will have an advantage.

Software:

- - SW providers should be part of EsD.
- - SW involved in EsD should be existing and running already before EsD starts.
- - Complementary H2020 calls for SW-development needed before the EsD.
- - No expected major SW development within EsD.
- - Both HW and SW components involved on the EsD should be "almost production ready".
- - Role of EsD is to integrate all components together.

SMEs:

- - Important for EsD: are more flexible and faster to engineer SW and HW than large companies.
- - Their task should be limited to very specific components (i.e. cooling, specific SW-layer or tool, etc.)
- - Important to reduce their burden in participating on EC-projects (less reporting, less administration, etc)
- - Encourage consortium, and in particular integrator to reserve some % of funding for SME.
- .Include SME on supply chain from integrator.
- .This facilitates further productisation of the SME product.
- - Involve also user-SMEs (application engineering, etc.)
- .Maybe create an industrial panel for that.
-
- Productisation:

- - Exploit markets outside HPC (server market, etc.)
- - "Keep it simple": find the right balance between risk/innovation and conservative approaches.
- - National co-funding only possible/justifiable if a product comes out of the EsD.
- - Important to address IP issues very early
- .Include some kind of business plan comprising IP-regulations in proposal.
-
-

Use cases Enabled by EsDs

- Mark Bull
- Alan O'Cais

Round 1

Reporter:

- Projects should be application driven, not hardware driven.
- Need to build in application requirements from the start.
 - Science case should be part of EsD bid
- Application porting and optimisation work has to be funded either as part of EsD or from elsewhere
 - latter may cause synchronisation problems
 - how do we motivate app devs to target what might be one-off systems?
 - Is 20 Pflop/s enough of a carrot?
- How do we provide HW designers with application requirements early enough?
 - modelling/simulation?
- Projects should be application driven, not hardware driven.
- Need to build in application requirements from the start.
 - Science case should be part of EsD bid
- Application porting and optimisation work has to be funded either as part of EsD or from elsewhere
 - latter may cause synchronisation problems
 - how do we motivate app devs to target what might be one-off systems?
 - Is 20 Pflop/s enough of a carrot?
- How do we provide HW designers with application requirements early enough?
 - modelling/simulation?

Round 2

Reporter: Alan O'Cais

There was a lot of discussion and I won't be able to do justice to it, but the basic outcomes were:

- The metric which is used to evaluate the proposals and their ultimate success should be science driven and not based on flop counts (it is the science that needs to be "exascale")

- The metric should be based on an set of scientific "Grand Challenges" (**not** particular applications) that cover a spectrum of fields important to Europe
 - The Grand Challenges should be included in the call text
 - There should be about 10 challenges
 - It should be up to EsD's how many of the challenges they target (intelligent grouping is obviously a good idea)
 - The mandate for selecting the challenges should be given to a capable group (perhaps a combination of the CoEs and PRACE?)
- The overall aim should be to bring entire applications and workflows to "exascale"
 - "co-design" work is only likely on elements of the application/work-flow, these elements need to be well selected
- The EsD's should have a production phase which is open to the entire European research community
 - Application partners doing the co-design to ensure that the Grand Challenges can be addressed will obviously benefit the most from this production phase
 - The production period should be extended to 3+ years to allow scientists external to the consortia a real opportunity to generate science on the EsD's

What remained unclear to us, from the use case perspective, is what the role of application-oriented industry and the CoE's is to be? If their participation is to be ensured then an appropriate (funding) strategy needs to be included to allow them to manage the risk and the potential sunk costs of the valuable productive PMs that are in short supply for such small organisations.

Additional Material

Also, in one of the "Case Studies" sections we have "agreed" on an approach that did not find its way in the final presentation to the community. Namely:

1) Applications (people) will express their interest to participate in the EsD consortia, without committing up front exclusively to any of them.

- These expressions can come from the CoE or from any other institutes (in order not to disadvantage these that skipped one round of funding)

- The expressions will contain the following elements:

currently running work volume (eg. grid size, input data file size, etc.)

characterization of the workload (eg. Flops, memory requirement, I/O need, etc.)

scientific goals of the run on the EsD (i.e. Achieve 24 hours forecast in 30 mins with resolution 1.5 km)

- There should be several (i.e. 10 or 20) such applications which will be grouped according to the requirements. For example, there could be a group with embarrassingly parallel apps, I/O dominant apps, etc. We did not exclude any app characteristic in this discussion.

- a question remained if and how apps will be funded to do this "advertisement" work ahead of any EsD consortia incorporation/inclusion.

2) Hardware consortia will pick applications, such that at least several groups are covered. It is a matter of balance and future discussion how many groups of apps must be covered in each EsD design. It is viewed that choosing the right balance will allow a degree of specialization of the EsD, while clearly several groups must be covered to avoid a single characteristic (i.e. Flops) to be over emphasized.

- applications hope to be payed from inclusion into the consortia, but it is up to other section to discuss
- EsD will be designed with the goal to demonstrate "efficient" run of the picked applications. Efficient in comparison to commercial machine available at the same time that will run the same app (we did not agree on the amount of tuning - it is a matter of balance)
- the scientific goals of the apps must be met with the EsD run