



Extreme scale Demonstrators

- Concept with Latest Updates -

EsD Roundtable @ HPC Summit Week 2017
Barcelona, May 18th 2017

Thomas Eickermann
Marc Durantou



May 18, 2017



- “The “Extreme-Scale Demonstrators” (EsDs) 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;
It is a key step towards establishing European exascale capabilities and solutions.” *(From the ETP4HPC SRA, chapter 8 p.67)*
- EsD will fill critical gaps in the HPC H2020 programme:
 - Bring technologies from FET-HPC closer to commercialisation (TRL 7-8)
 - Combine results from targeted R&D efforts into a complete system (European HPC technology ecosystem)
 - Provide the missing link between the 3 HPC pillars: Technology providers, infrastructure providers, user communities (co-design)

Technology providers

- Technology integration
- System architects
- Testing and quality/performance assurance (phase A)
- Maintenance and service (phase B)

EsD Expectations

- Design points ~400-500 Pflops
- EsD target 5% (20-30 Pflops)
- Budget: 20-50 Mio. €
- Diversity of architectures
- TRL 7-8

EsDs

Application owners / CoEs

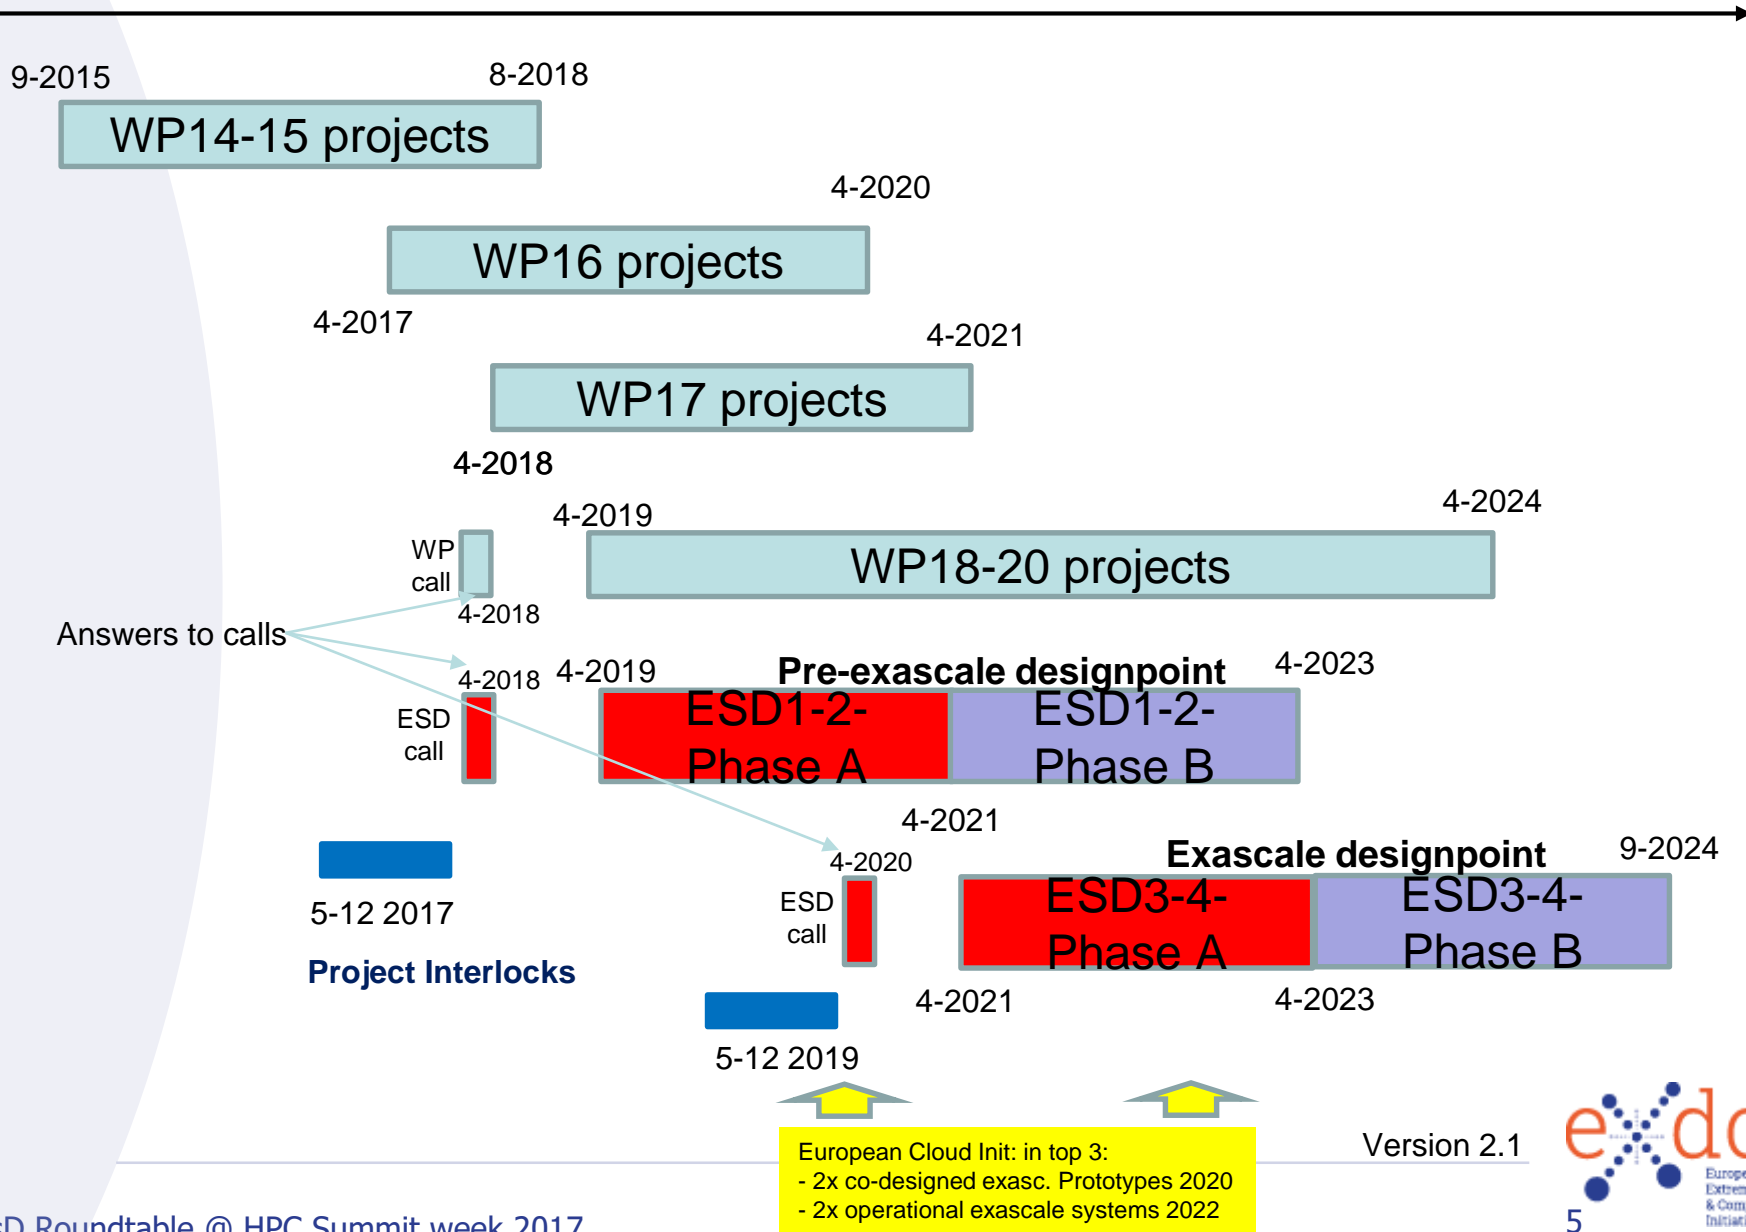
- Application requirements and key challenges (phase A)
- Port, optimize application(s), use them productively (phase B)

HPC Centres

- Participate in co-design
- Manage system deployment (phase A)
- System operation, validation (phase B)

- Two EsD calls proposed by ETP, each leading to two projects
 - Calls target technologies developed under FETHPC, but are open
 - EsD projects should start after end of FETHPC projects in WP2014/15 and WP2016/17
- EsD project structure
 - Phase A (18-24 months): Development, Integration and Testing
 - Little or no basic technology research
 - Substantial R&D focus geared towards integrating components and sub-systems developed in the preceding R&D projects
 - Phase B (18-24 months): Deployment and Use
 - Operated by a hosting center
 - EsD made available to application owners for code porting and development
 - Characterization and EsD validation, benchmarking based on real use-cases.

Time Line of the HPC cPPP (Slightly Outdated)



EsD State of Play (#1/2)

- EsD concept is outlined in the latest ETP4HPC SRA (to be detailed in SRA3)
- Broad discussion on EsDs in several Workshops:
 - EXDCI Workshop '15, Rome: Kick-Off
 - HPC Summit '16, Prague: Discuss EsD concept in community and with EC:
 - Technical system characteristics, budget, procurement model, consortia, use cases
 - ISC '16, Frankfurt:
 - continue Prague discussion
 - focus: from Use Cases to System Specs
 - Results available at <http://www.etp4hpc.eu/en/esds.html>
The Extreme-Scale Demonstrators concept: Current State of Definition
 - EXDCI Workshop '16, Barcelona:
 - focus: system integrators & Application Communities
 - gather input on expectations, interests, constraints, ...

EsD State of Play (#2/2)

- Interaction with EC
 - Various informal discussions with EC
 - cPPP Board meeting, Nov. '16, Brussels:
 - Present status to ETP and EC
 - Pose and discuss open questions / issues
 - Conclusion: focused workshop with EC needed
(in particular on budget, scope of calls, acquisition model)
 - EsD Workshop with EC, Dec. 5th '16, Brussels: some clarifications / **see next slides**
- ETP has submitted recommendations for WP18-20 in December 2016
 - EC's feedback received and discussed, reply provided / **see following slides**
(affects: number of calls, overall budget, performance targets)
 - No further input from ETP to WP18-20 is foreseen
- HPC Summit Week, May 18th '17, Barcelona
 - EsD Roundtable: final exchange before WP is published & call is opened

EC feedback: Call Design & Funding Model

- No special treatment for EsD – operate within H2020 rules
- Call design
 - 2 separate calls require separate sets of objectives
 - Reference to FETHPC in WP14-15 vs. WP16-17 is not sufficient
 - EsDs of first call should contribute to pre-exascale systems
 - 1-stage evaluation (2-stage is rarely used)
 - Cascade funding is too complex; rely on Contract Amendments
- Funding envelope and funding model
 - How much funding can reasonably be absorbed at a given point in time?
 - Only RIA offers 100% funding rate → *requires R&D in Phase A*
- Acquisition model
 - PCP and PPI provide < 100% funding rate
 - Usual rules for eligible costs apply: staff, other, subcontracting
 - Integrator should be part of the consortium
→ *Purchase by system integrator; rules do not allow profits*

ETP4HPC Recommendations for WP18-20

- General concept as described above
- Challenge
 - Complete hardware and software systems
 - Developed in co-design, usable in production-like mode
 - R&D focus: integrate, customise results of previous R&D projects + market
 - FP7, H2020, other R&D in Europe
 - Demonstrate scalability to Exascale, to become the basis of European capability in Exascale
 - Progress in energy-efficiency and cost of ownership
- Scope
 - Target 5% of performance of a high-end production system in a relevant metric (*Flops not mentioned!*)
 - Applications: port, adapt, develop towards new programming models, algorithmic approaches

ETP4HPC Recommendations for WP18-20

- Scope of Call 1 (2018)
 - WP14-15, FP7
 - Demonstrate scalability to Exascale (e.g. 0,5 ... 1 EFlops)
 - Target commercially viable, competitive Exascale in 2021/2022
 - Demonstrate operational and performance for a number of key applications
- Scope of Call 2 (2020)
 - WP16-17
 - Target new application areas, e.g. HPDA, Machine Learning, ...
 - Extreme computing valid, if sufficient progress since Call 1
 - Demonstrate operational and performance for a number of new application areas/use cases
 - key applications
- Funding and Duration
 - 100 Mio € / call, 2 projects / call
 - Phase A ~24 months, Phase B ~ 24 months (extended operation appreciated)

EC Feedback on ETP4HPC Recomm. & ETP Comments

- Relation of EsDs to the European low power CPU and contribution to Exascale machines with European technology to be more clearly outlined
 - ETP did not have sufficient and timely information to address this in its recommendations
 - ETP supports within European HPC programmes the use of new low power processors that will enable worldwide competitive European HPC system designs
- Can the 2nd EsD call and the "Major system level co-design projects" in the RIAs be replaced by an integrated/closer related effort?
 - These actions have different scope and objectives and should be kept separate
 - 2nd EsD call: take-up of WP16-17 results and new applications areas, bringing it close to commercialization
 - Major system level co-design projects: continuation of the bigger co-design projects started in 2016
- Provide other scalability/performance design points than Flops
 - Better not bias and limit the design space (e.g. Bytes/Flops, W/Flops, Flops/rack suitable for today's number crunchers)

Conclusions

- Concept of EsDs welcomed by the EC
- Expected to materialise in the upcoming WP18-20
- Characteristics may differ from ETP recommendations in
 - Number and timing of calls
 - Funding and expected number of projects
 - Specific objectives (performance characteristics, targets)
- Apparently different expectations from EsDs:
 - EC: primarily a milestone towards the European Technology based Exascale Systems
 - ETP: continuous effort required to bring results from RIA projects closer to commercialisation

Thank you for your Attention !

Questions ?

Backup Slides

EC Feedback on ETP4HPC Recomm. & ETP Comments

- Relation of EsDs to the European low power CPU and contribution to Exascale machines with European technology to be more clearly outlined
 - ETP was not aware of European Low power microprocessor calls.
 - Without notion of the scope, the intended characteristics in a multiyear roadmap, the targeted markets, the competitive positioning and differentiation of such a chipset, ETP4HPC is not in a position to consider it in the WP18-20 recommendations.
 - ETP supports within European HPC programmes the use of new low power processors that will enable worldwide competitive European HPC system designs.

EC Feedback on ETP4HPC Recomm. & ETP Comments

- Can the second Call for Extreme Scale Demonstrators and the "Major system level co-design projects" in the Research and innovation actions be replaced by an integrated/closer related effort?
 - These actions have different scope and objectives and should be kept separate
 - 2nd EsD call targets take-up of Wp16-17 results and new applications areas, bringing it close to commercialization
 - "Major system level co-design projects" is intended as a continuation of the bigger co-design projects started in 2016

EC Feedback on ETP4HPC Recomm. & ETP Comments

- There is only one criterion for EsDs 2018 (on scalability/performance design point), which is insufficient. Other criteria should be identified (e.g. power consumption), and whenever possible, giving concrete targets.
 - We did not put up more criteria to not bias and limit the design space, on purpose.
 - As an alternative a range of indicative figures could be mentioned taken partly from the SRA:
 - Scalability design point: 500 PFLOPs to 1 EFLOPS
 - Balanced architecture : in particular wrt memory BW. i.e. no more than 2 Bytes/Flops
Energy efficiency :35kW/PFLOPS
 - Compute packaging density: 1 PFLOPS/rack
 - Total number of racks per system:4-6

Conclusions from Barcelona Workshop

- Integrators presented their view
 - Atos, Cray, E4, EuroTech, Fujitsu, Huawei, Lenovo, Megware
- Preparation of a proposal requires several months (time & effort)
 - Integration of technologies from different FETHPC projects
 - Design to a level that allows a cost estimate
 - → Workshop with FETHPC projects in spring 2017
 - Other supporting measures?
- Critical factors for vendors to participate:
 - Funding and acquisition model (to be discussed with EC)
 - Management of IP in consortia (to be solved by consortia)
 - EsDs must not be one-of systems, but open a product line (first-of): Substantial investment of project partners needs to pay off
- Application communities ask for true co-design cycle
 - Users provide requirements → architects design system will not work

Options for Acquisition

- PCP + PPI
 - + Well-defined, innovation-driven process
 - Long process: 3+ years
 - PPI requires significant contribution from procurer
 - Strict separation between procurer and supplier
- Public procurement by hosting partner within RIA project
 - + Well established process; HPC centres have strong experience
 - Public procurement rules may conflict with call objectives
 - Need to avoid open procurement; other providers may challenge it ...
- Purchase by system integrator within RIA
 - + Simple, well established process
 - Integrator cannot buy from itself (need to purchase from manufacturers)
 - Financial risks may be too high for smaller companies – excludes SME