



EXDCI workshop

Barcelona, Sept. 21 and 22, 2016

Michael Malms



September 22, 2016



BDVA - use case discussion

EXDCI workshop
Barcelona, Sept. 21st 2016

Jim Kenneally

Michael Malms

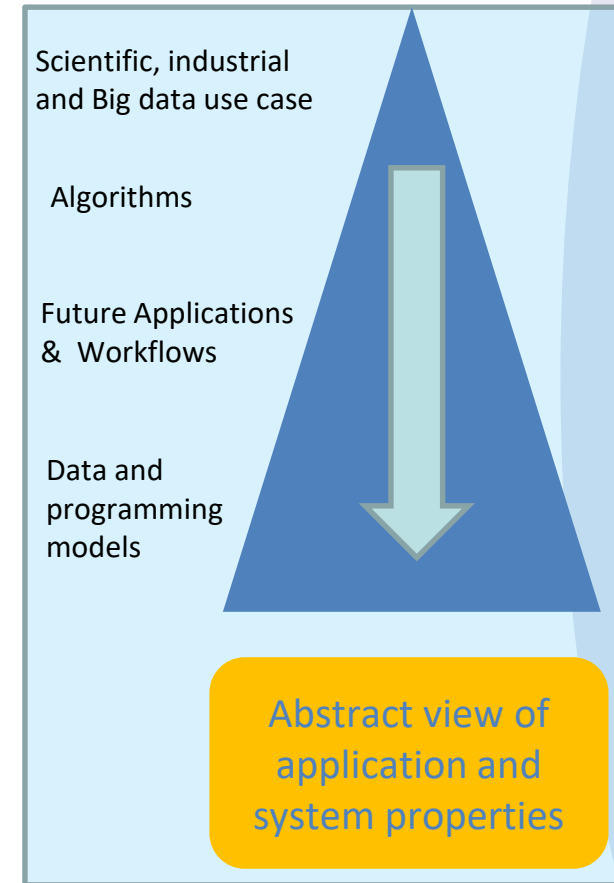


September 22, 2016



Some words upfront....

- In WP2 of EXDCI, a new release of the Strategic Research Agenda (SRA) is due in 7/2017
- The use of heavy Big Data / HPDA applications will increasingly take advantage of HPC compute infrastructures
- Therefore, future HPC system architectures will have to accommodate specific BD/HPDA architectural requirements.
- In order to understand the drivers of the influences, a top-down approach seems most suited
- In this session BD use cases will be reviewed with the intention to start deriving an abstracted view of application and system properties:
 1. Healthcare
 2. Transport
 3. Civil Safety
 4. Natural Language Processing
 5. Performance aware Big Data processing



WP2-WP3 interlock

EXDCI workshop
Barcelona, Sept. 21st 2016

Stephane Requena
Michael Malms

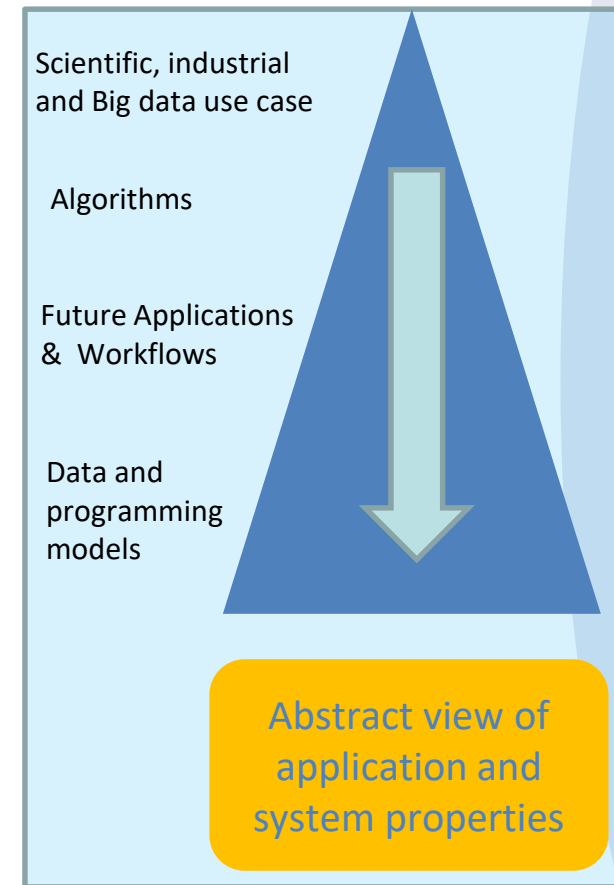


September 22, 2016



Some words upfront....

- In WP2 of EXDCI, a new release of the Strategic Research Agenda (SRA) is due in 7/2017
- Today, scientific and industrial use cases in the domain of technical computing still dominate the architecture of the HPC compute infrastructure
- On the path to exascale it is important to understand the 2022/23 workload demands in this application domain
- In order to understand the drivers of the influences, a top-down approach seems most suited
- In this session use cases in recognized scientific HPC application domains will be reviewed with the intention to start deriving an abstracted view of application and system properties:
 1. Domain 1
 2. Domain 2
 3. Domain 3
 4. Domain 4



Extreme scale Demonstrators: some words upfront (2)

It is important for the definition of the next layer of the EsD concept to understand and discuss the HPC infrastructure / related requirements of major HPC scientific and industrial application domains.

Suggested topics for the presentations of the WP3 scientific & industrial use cases:

1. What are the scientific problems to be solved ?
2. What is the tangible goal for 2022/23 ?
3. What is hampering you from achieving this goal today? List all technical and operational roadblocks / hurdles.
4. If you have specific HPC infrastructure requirements and recommendations for the 2022/23 timeframe, what are they ?



Extreme scale Demonstrators

- the voice of system integrators -

EXDCI workshop
Barcelona, Sept. 21st 2016

Thomas Eickermann
Michael Malms



September 22, 2016



Extreme scale Demonstrators: some words upfront (1)

- Agenda for this session:
 - Quick recap of the current status of the EsD proposal (Thomas Eickermann, 20 min)
 - System integration related discussion points (next page) (Michael Malms, 10 min)
 - Presentation of System Integrators (10 min each)
 - Lenovo
 - Cray
 - Megware
 - e4
 - Atos/Bull
 - Eurotech
 - Fujitsu
 - Huawei
 - Discussion „essential steps towards executable EsD projects“ (All, 60 min)

Extreme scale Demonstrators: some words upfront (2)

System integrators play a pivotal role in the EsD concept both during Phase „A“ (development & integration) as well during phase „B“ (evaluation/ benchmark and deployment).

(Note: the term „system integrator“ is vague and can span a variety of roles and steps towards making a system „shippable“).

It is important for the definition of the next layer of the EsD concept to understand and discuss the different positions of interested EsD system integrators such as:

Suggested topics for the presentations of the system integrators:

1. What is your motivation for potentially assuming the role of a system integrator in one of the Esd projects ?
2. What are the responsibilities and scopes you wish to cover:
 - System architect
 - Development of own subsystem(s) or subcomponent(s)
 - Integration of entire system including third party subsystems & components
 - System Test and EsD-release
 - Maintenance & support
3. What is your view on how best to implement the object of „integrate technology researched and prototyped in previous FETHPC projects“? (analysis-process, estimated time and resources required, IP-aspects, etc.)
4. What is your view on the governance – structure of an EsD project, what are your fundamental requirements?
5. Any thoughts about required budgets and funding mechanisms ?
6. Any „no go“ aspects ?



Strategic Research Agenda SRA

a multi-annual roadmap towards
Exascale High-Performance Computing
Capabilities

September 22, 2016



ETP4HPC

European Technology Platform
for High-Performance Computing

Strategic Research Agenda 2015 Update

European Technology
Multi-annual Roadmap
Towards Exascale
Update to 2013 Roadmap



H2020-FETHPC-2014

Coordination of the HPC strategy



EXDCI

European eXtreme Data and Computing Initiative

Grant Agreement Number: FETHPC-671558

D2.1

Update of Strategic Research Agenda (SRA2)

Final

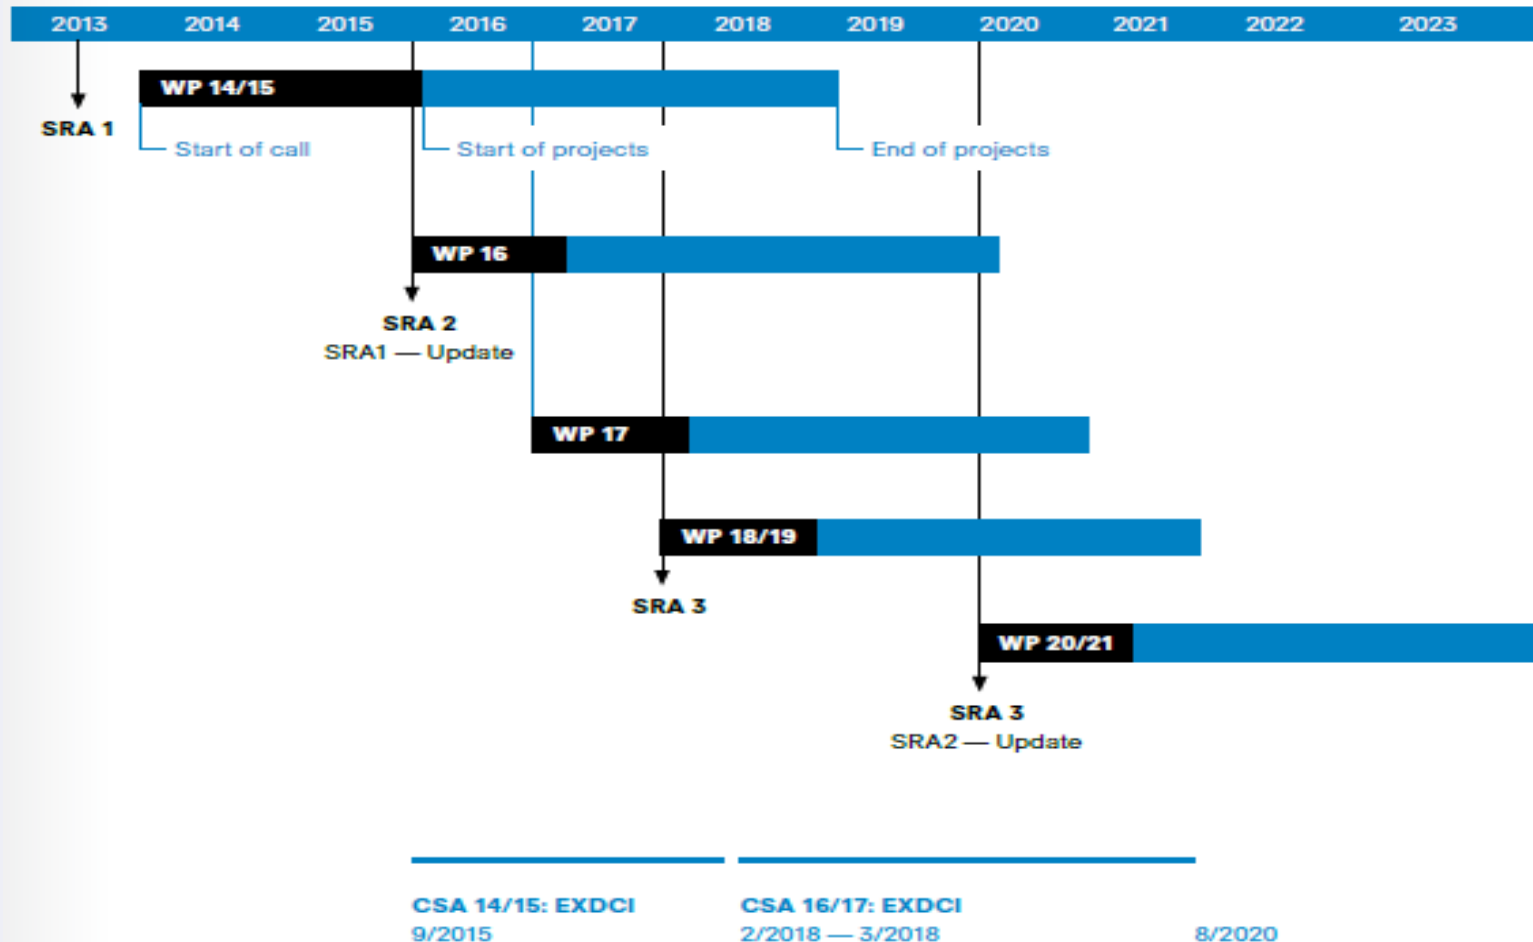
Version: 1.1
Author(s): Michael Malm, Jean-Philippe Nominé, Marcin Ostasz (ETP4HPC)
Date: 19.01.2016



September 22,
2016

Horizon 2020 WPs and SRAs

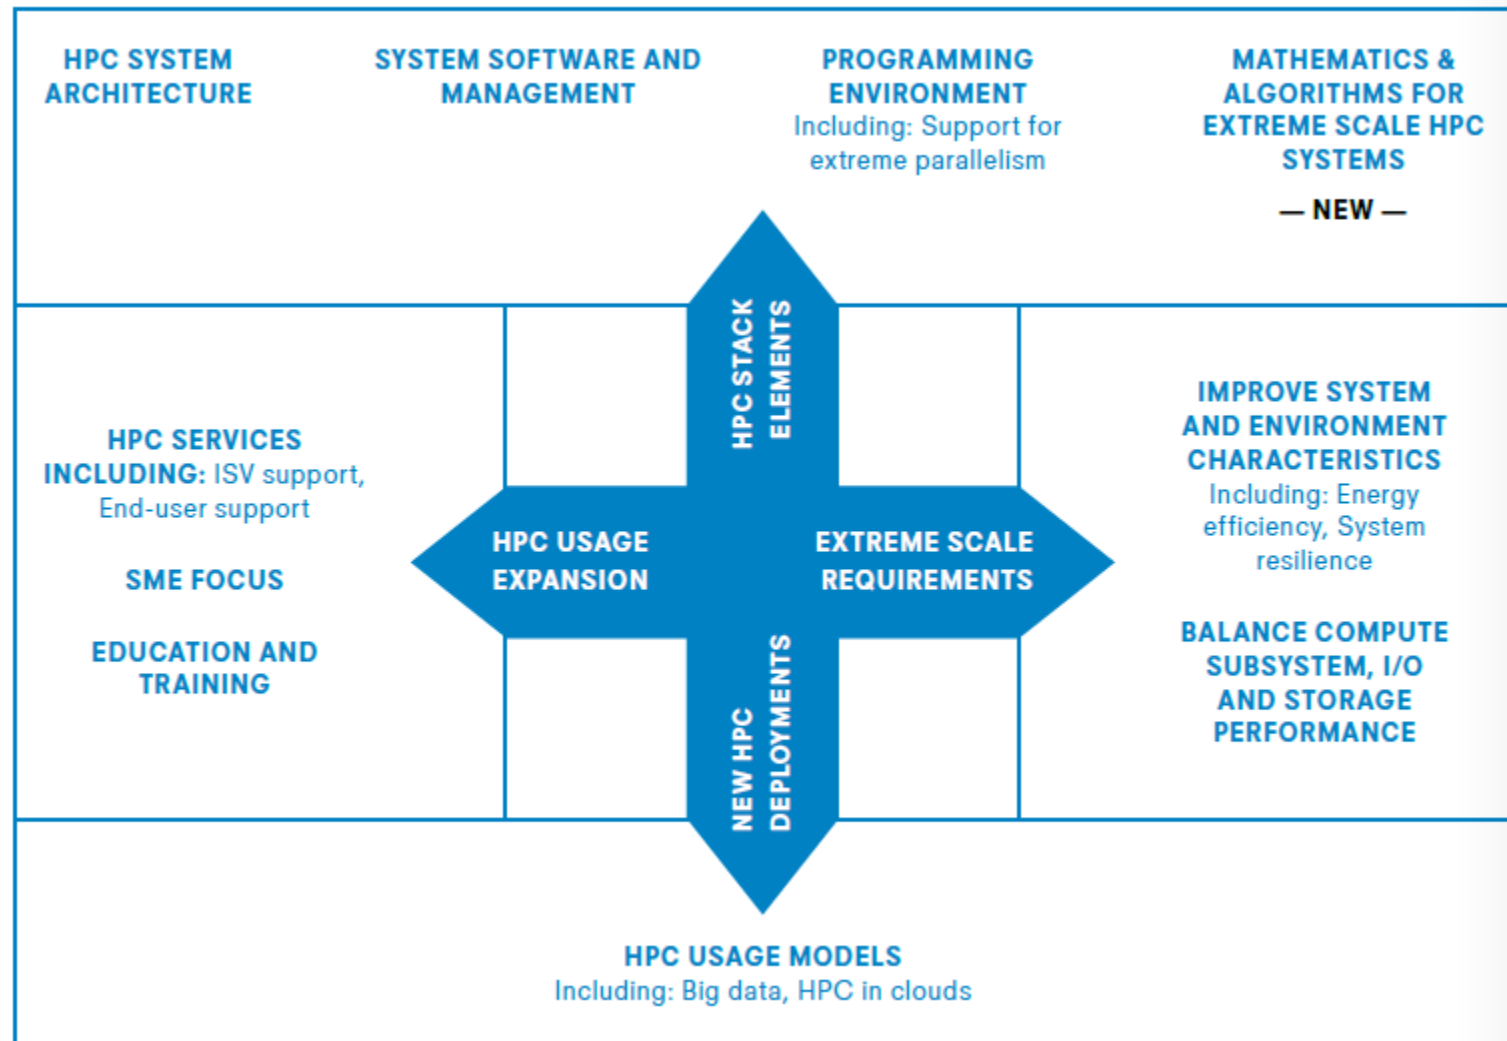
HPC — HORIZON 2020 ROADMAP



Priorities

- There is a demand for R&D and innovation in both extreme performance systems and mid-range HPC systems
 - Scientific domain and some industrial users want extreme scale
 - ISVs and part of the industry expect more usability and affordability of mid-range system
- The ETP4HPC HPC technology providers are also convinced that to build a sustainable ecosystem,
 - their R&D investments should target not only the exascale objective (too narrow a market)
 - an approach that aims at developing technologies capable of serving both the extreme-scale requirements and mid-market needs can be successful in strengthening Europe's position.

4 dimensions of the SRA



Transversal issues to be addressed

- Three technical topics:
 - Security in HPC infrastructures to support increasing deployment of HPDA
 - Resource virtualisation to increase flexibility and robustness
 - HPC in clouds to facilitate ease of access
- Two key elements for HPC expansion
 - Usability at growing scale and complexity
 - Affordability (focus on TCO)

How has the SRA been built?

- 8 Workgroups covering the 8 technical focus areas:

SRA 2015 technical focus areas

- HPC System Architecture and Components
 - Energy and Resiliency
 - Programming Environment
 - System Software and Management
 - Big Data and HPC usage Models
 - Balance Compute, I/O and Storage Performance
 - Mathematics and algorithms for extreme scale HPC systems
 - Extreme scale demonstrators
-
- 48 ETP4HPC member orgs/companies involved in these workgroups
 - Members named 170 individual experts to contribute, 20-30 per working group

Other interactions

- Feedback sessions with end-users and ISVs at Teratec Forum
 - 20 end-users outline their deployment of HPC, future plans and technical recommendations
 - Very diverse set of priorities (performance & scale, robustness, ease of access, new workflows etc.)
 - No 'One size fits all' – approach possible
- Technical session with Big Data Value Association (BDVA) to understand architectural influences of HPDA
 - Technical dialogue started, much more to be done over next 1-2 years
 - BDVA has issued an update to their SRIA in Jan 2016

The technical domains and the ESD proposal

Trends and recommended research topics – a few examples

September 22, 2016



HPC System Architecture, Storage and I/O, Energy and Resiliency

Major trends - a subset:

- Increased use of accelerators (e.g. GPUs, many core CPUs) in heterogeneous system architectures
- Compute node architectures efficiently integrate accelerators, CPUs with high bandwidth memory
- Non volatile memory types open up new interesting memory and caching hierarchy designs
- System networks to significantly scale up and cut latencies, introducing virtualisation mechanisms
- Storage subsystems to become more 'intelligent' to better balance compute and I/O
- Increased activities in object storage technologies with major architectural revamp in the next years
- Focus on architectural changes to improve energy efficiency and reduce data movement

Research topics to be addressed (examples)

- Compute node deep integration with embedded fast memory and memory coherent interfaces
- Silicon photonics and photonic switching in HPC system networks
- Global energy efficiency increases with targets of 60kW/PFlops in 2018 and 35 kW in 2020
- Active storage technologies to enable 'in situ' and 'on the fly' data processing
- Research in methods to manage 'energy to solution'
- Prediction of failures and fault prediction algorithms

HPC System Architecture, Storage and I/O: milestones

M-ARCH-1: New HPC processing units enable wide-range of HPC applications.	2018
M-ARCH-2: Faster memory integrated with HPC processors.	2018
M-ARCH-3: New compute nodes and storage architecture use NVRAM.	2017
M-ARCH-4: Faster network components with 2x signalling rate (rel. to 2015) and lower latency available.	2018
M-ARCH-5: HPC networks efficiency improved.	2018
M-ARCH-6: New programming languages support in place.	2018
M-ARCH-7: Exascale system energy efficiency goals (35kW/PFlops in 2020 or 20 kW/Pflops in 2023) reached.	2020-2023
M-ARCH-8: Virtualisation at all levels of HPC systems.	2018
M-ARCH-10: New components / disruptive architectures for HPC available.	2019

M-BIO-1: Tightly coupled Storage Class Memory IO systems demo.	2017
M-BIO-2: Common I/O system simulation framework established.	2017
M-BIO-3: Multi-tiered heterogeneous storage system demo.	2018
M-BIO-4: Advanced IO API released: optimised for multi-tier IO and object storage.	2018
M-BIO-5: Big Data analytics tools developed for HPC use.	2018
M-BIO-6: 'Active Storage' capability demonstrated.	2018
M-BIO-7: I/O quality-of-Service capability.	2019
M-BIO-8: Extreme scale multi-tier data management tools available.	2019
M-BIO-9: Meta-Data + Quality of Service exascale file i/o demo.	2020
M-BIO-10: IO system resiliency proven for exascale capable systems.	2021

Energy and resiliency: milestones

M-ENR-MS-1: Quantification of computational advance and energy spent on it.	2017
M-ENR-MS-2: Methods to steer the energy spent.	2017
M-ENR-MS-3: Use of idle time to increase efficiency.	2018
M-ENR-AR-4: New levels of memory hierarchy to increase resiliency of computation.	2017
M-ENR-FT-5: Collection and Analysis of statistics related to failures.	2018
M-ENR-FT-6: Prediction of failures and fault prediction algorithms.	2019

M-ENR-FT-10: Application survival on unreliable hardware.	2019
M-ENR-AR-7: Quantification of savings from trade between energy and accuracy.	2018
M-ENR-AR-8: Power efficient numerical libraries.	2019
M-ENR-MS-9: Demonstration of a sizable HPC installation with explicit efficiency targets.	2019

Extreme-Scale Demonstrators

- **Characteristics**

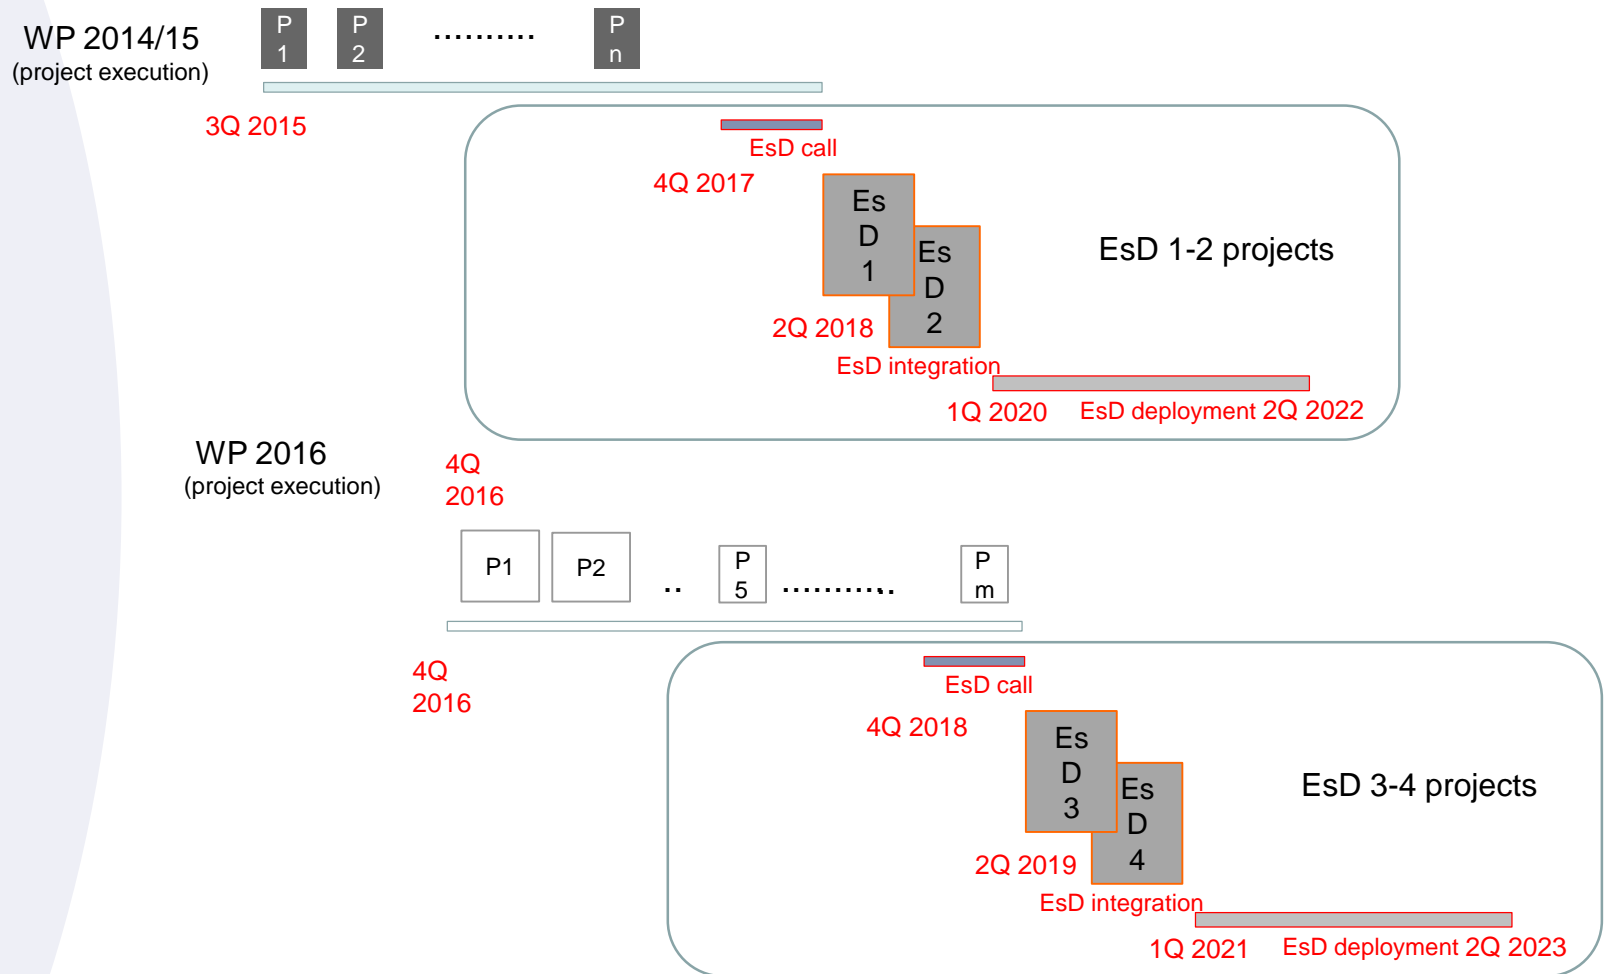
- Four complete prototype HPC systems, calls in 2018 & 2019
- high enough TRL to support stable production
- using technologies developed in the previous projects
- based on application – system co-design approach
- large enough to address scalability issues (at least 5% of top performance systems at that time)

- **Two project phases:**

- phase A : development, integration (of results from R&D projects) and testing
- phase B : deployment and use, code optimisation, assessment of the new technologies

Extreme scale Demonstrators call-integration-deployment

schedule



SRA – next actions

September 22, 2016





We will welcome your comments on the current SRA
<http://www.etp4hpc.eu/strategic-research-agenda/>

Strategic Research Agenda | ETP4HPC

www.etp4hpc.eu/strategic-research-agenda/ ▼

6 days ago - Public Call for Comments on ETP4HPC Strategic Research Agenda. Our organisation would like to receive feedback on this document from the ...

Public Call for Comments for ETP4HPC Strategic Research ...

<https://www.surveymonkey.com/.../ETP4HPC-SRA2-PUBLIC-CALL4C...> ▼

The updated Strategic Research Agenda (SRA) of ETP4HPC is now available at the following location: <http://www.etp4hpc.eu/strategic-research-agenda/>

Public Call for Comments on ETP4HPC Strategic Research ...

primeurmagazine.com/flash/AE-PF-12-15-16.html ▼

2 days ago - Public Call for Comments on ETP4HPC Strategic Research Agenda for exascale supercomputing in Europe December 2015. 13 Dec 2015 ...

Primeurflash 20151213 - Primeur Magazine

primeurmagazine.com/contentsflash20151213.html ▼

2 days ago - Public Call for Comments on ETP4HPC Strategic Research Agenda for ... Agenda on November 24th 2015, the ETP4HPC organisation would ...

ETP4HPC, EXDCI and SESAME Net - new HPC initiatives in ...

e-irg.eu/.../etp4hpc-exdci-and-sesame-net-new-hpc-initiatives-in-europe-... ▼

Apr 9, 2015 - The HPC Centres of Excellence Call amounts to 14 million euro. ... will require an investment of 15 million euro; the Public Procurement of innovative HPC systems has been estimated at 26 million; 698 Views, 0 Comments.
 You visited this page on 12/2/15.

Catherine Gleeson | LinkedIn

<https://www.linkedin.com/in/catherine-gleeson-151229b7>

Amsterdam Area, Netherlands - ETP4HPC - European Technology Platform for HPC - ETP4HPC

Catherine Gleeson. ETP4HPC - European Technology Platform for HPC ... Public Call for Comments on ETP4HPC Strategic Research Agenda. December 11 ...

eInfrastructures (@eInfraEU) | Twitter

<https://twitter.com/einfraeu> ▼

"Public Call for Comments on ETP4HPC Strategic Research Agenda" by @Etp4H on ... New #H2020 #einfrastructures call for support to policy and international ...

Images for etp4hpc public call for comments

Report images



Next SRA-related events in 2016

- HPC summit – Extreme scale Demonstrator workshop - May 12th
 - focussed on the EsD definition (engage potential players, further implementation details)
 - at this event the three pillars for the EsD mission (CoE, HPC centres and the FETHPC1 project speakers) are invited . **More than 80 registered participants!**
- Participation in BDEC conference - June 16 & 17
- ISC16 – June 23rd
 - Scope: Feedback session on SRA directions, content and value to shape the next update
(Invited are: End-users, ISVs and International HPC experts)
 - 2nd EsD workshop (follow-on to May 12th workshop)
- Level set with HPC application experts (EXDCI WP3) – September 21 & 22
- Technical workshop with Big Data Value Association (BDVA) June/July



THANK YOU!

For more information visit

www.etp4hpc.eu

contact: office@etp4hpc.eu

September 22, 2016



Backup

September 22, 2016



Major trends – a subset:

- New node architectures demand innovative methods to solve scalability and concurrency issues
- Network virtualisation and data security become critical system level challenges
- Support for increasing use of 'in situ' data processing
- Driven by HPDA, resource management needs to cope with highest levels of data allocation flexibility
- Increased intelligence throughout the programming workflow
- Productivity enhancements through use of domain specific languages (DSLs)
- Interoperability and composability of programming models provide more flexibility to appl. developer

Research topics to be addressed (examples)

- Efficient OS support for heterogeneous architectures with complex memory hierarchies
- Congestion control and adaptive /dynamic routing algorithms for exascale interconnects
- Research on data-aware scheduling and resource management
- Programming tool intelligence based on cost models for e.g. energy used, load-balancing, etc.
- Programming models to allow for malleability (ability to adapt to changing resource availability)

System Software and Management: milestones

M-SYS-OS-1: Kernel scheduling policy.	2016
M-SYS-OS-2: OS Low level standard API with run-time.	2017
M-SYS-OS-3: New memory management policy and libraries.	2017
M-SYS-OS-4: Container and virtualisation support; Hypervisor for HPC.	2016
M-SYS-OS-5: Offload programming model support.	2017-2019
M-SYS-OS-6: OS decomposition to add application performance and flexibility.	2019
M-SYS-OS-7: Investigate HPC specific security requirements on OS level.	2017-2019
M-SYS-IC-1: OS-bypass and hardware interface integrity protection.	2016
M-SYS-IC-2: Interconnect adaptive and dynamic routing algorithm and congestion control, power management.	2017
M-SYS-IC-3: Network virtualisation compliancy.	2017

M-SYS-CL-1: Flexible execution context configuration and management (from image to containers).	2018
M-SYS-CL-2: Prescriptive maintenance based on Big Data analytics technics.	2016
M-SYS-CL-3: Infrastructure security.	2017-2020
M-SYS-RM-1: New Scalable scheduling enhancement, with execution environment and data provisioning integration.	2017
M-SYS-RM-2: New multi-criteria adaptive algorithms: Heterogeneity-/memory- and locality-aware.	2017
M-SYS-RM-3: Resilient framework.	2020
M-SYS-Vis-1: Scalable "in situ" visualisation.	2016
M-SYS-Vis-2: Scaling for the compositing phase.	2017
M-SYS-Vis-3: Ray-tracing capabilities.	2018
M-SYS-Vis-4: High dimensional data, graphs and other complex data topologies.	2018

Programming Environment: milestones

M-PROG-API-1: Develop benchmarks and mini-apps for new programming models/languages.	2016
M-PROG-API-2: APIs and annotations for legacy codes.	2017
M-PROG-API-3: Advancement of MPI+X approaches (beyond current realisations).	2017
M-PROG-API-4: APIs for auto-tuning performance or energy.	2017
M-PROG-API-5: Domain-specific languages (specific languages and development frameworks).	2018
M-PROG-API-6: Efficient and standard implementation of PGAS.	2018
M-PROG-API-7: Non-conventional parallel programming approaches (i.e. not MPI, not OpenMP / pthread / PGAS - but targeting asynchronous models, data flow, functional programming, model based).	2019
M-PROG-LIB-1: Self- / auto-tuning libraries and components.	2018
M-PROG-LIB-2: Components / library interoperability APIs.	2017
M-PROG-LIB-3: Templates / skeleton / component based approaches and languages.	2019
M-PROG-RT-1: Run-time and compiler support for auto-tuning and self-adapting systems.	2018
M-PROG-RT-2: Management and monitoring of run-time systems in dynamic environments.	2018

M-PROG-RT-3: Run-time support for communication optimisation and data placement: data locality management, caching, and prefetching.	2019
M-PROG-RT-4: Enhanced interaction between run-time and OS or VM monitor (w.r.t. current practice).	2018
M-PROG-RT-5: Scalable scheduling of million-way multi-threading.	2020
M-PROG-DC-1: Data race condition detection tools with user-support for problem resolution.	2017
M-PROG-DC-2: Debugger tool performance and overheads (in CPU and memory) optimised to allow scaling of code debugging at peta- and exascale	2018
M-PROG-DC-3: Techniques for automated support for debugging (static, dynamic, hybrid) and anomaly detection, and also, for the checking of programming model assumptions.	2018
M-PROG-DC-4: Co-design of debugging and programming APIs to allow debugging to be presented in the application developers original code, and also, to support applications developed through high-level model descriptions.	2018
M-PROG-PT-1: Scalable trace collection and storage: sampling and folding.	2018
M-PROG-PT-2: Performance tools using programming model abstractions.	2018
M-PROG-PT-4: Performance analytics tools.	2018
M-PROG-PT-5: Performance analytics at extreme scale.	2019

- Major trends – a subset:

- Data analytics, including visualisation increasingly will take place 'in situ'
- HPC systems with lots of memory and fast networks become ideal compute infrastructure for Big Data
- Focus on math and algorithms for exascale system software (compilers, libraries, programming environment)
- Advances in mathematical methods req. to improve energy efficiency by two orders of magnitude

- Research topics to be addressed (examples)

- Research on new performance metrics to reflect data-centric use of HPC infrastructure
- Data centric memory hierarchies and architectures, data structure transformation to enable HPDA
- Systematic analysis of data flows in key Big Data applications to minimise data access and movement
- Research on HPC and Big Data hybrids to allow simulation and data analytics at the same time
- Mathematical support for data placement and data movement minimization
- Research on the impact of algorithmic and mathematical advances to programming tools
- Work on new algorithms to reduce energy to solution

Big Data and HPC Usage Models, mathematics and algorithms

M-BDUM-METRICS-1: Data movement aware performance metrics.	2017
--	------

M-BDUM-METRICS-2: HPC like performance metrics for Big Data systems.	2017
--	------

M-BDUM-METRICS-3: HPC-Big Data combined performance metrics.	2018
--	------

M-BDUM-MEM-1: Holistic HPC-Big Data memory models.	2017
--	------

M-BDUM-MEM-2: NVM-HPC memory and Big Data coherence protocols and APIs.	2017
---	------

M-BDUM-ALGS-1: Berkeley Dwarfs determination for Big Data applications.	2017
---	------

M-BDUM-ALGS-2: Implementations of Dwarfs in Big Data platforms.	2019
---	------

M-BDUM-PROG-1: Hybrid programming paradigms HPC-Big Data.	2017
---	------

M-BDUM-PROG-2: Hybrid programming paradigm with coherent memory and compute unified with Big Data programming environments.	2018
---	------

M-BDUM-PROG-3: Single programming paradigm across a hybrid HPC-Big Data system.	2021
---	------

M-BDUM-VIRT-1: Elastic HPC deployment.	2018
--	------

M-BDUM-VIRT-2: Full virtualisation of HPC usage.	2021
--	------

M-BDUM-DIFFUSIVE-1: Big Data - HPC hybrid prototype.	2017
--	------

M-BDUM-DIFFUSIVE-2: Big Data - HPC large-scale demonstrator.	2020
--	------

Mathematics and algorithms for extreme scale HPC systems: milestones

M-ALG-1: Scalability of algorithms demonstrated for forward in time computing for current architectures.	2017
M-ALG-2: Multiple relevant use cases demonstrated for improving performance by means of robust, inexact algorithms.	2018
M-ALG-3: Scalable algorithms demonstrated for graph-based analytics.	2019
M-ALG-4: Processes established for co-design of mathematical methods for data analytics and of HPC technologies/architectures.	2019
M-ALG-5: Classes of data, partitioning and scheduling problems categorised and their complexity ascertained.	2019
M-ALG-6: Mathematical and algorithmic approaches established for the scheduling of tasks on abstract resources and exploitation of multiple memory levels.	2020
M-ALG-7: Research on mathematical methods and algorithms exploited for compiler technologies, run-time environments and related tools.	2018
M-ALG-8: Reduction of energy-to-solution demonstrated by means of appropriately optimized algorithms demonstrated for a set of relevant use cases.	2017
M-ALG-9: Process for vertical integration of algorithms established together with the validation of scalability, ease of implementation, tuning and optimisation.	2019