

European Technology Platform for High Performance Computing



Vision Paper

EXECUTIVE SUMMARY

High Performance Computing (HPC) plays a pivotal role in stimulating Europe's economic growth. HPC is a pervasive tool allowing industry and academia to develop world class products, services and inventions in order to maintain and reinforce Europe's position on the competitive worldwide arena. HPC is also recognized as crucial in addressing grand societal challenges. "Today, to Out-Compute is to Out-Compete" best describes the role of HPC.

To sustain its competitiveness, Europe needs to enhance its innovation capabilities using HPC. However:

- Today, HPC is mostly used in large HPC research centres and some large enterprises. It needs to become much more pervasive across the entire industry, including Small and Medium Enterprises (SMEs).
- While Europe represents a significant portion of the world's available HPC resources, only a fraction of that technology and infrastructure is developed in Europe.
- Other regions invest considerably more in new technologies, architectures and software for HPC systems, gaining a substantial competitive edge over Europe.

An industry-led initiative, a European Technology Platform (ETP) for High Performance Computing, is being formed to address these issues. This ETP will define research priorities for the development of a globally competitive HPC technology ecosystem in Europe. It will propose and help to implement a Strategic Research Agenda, while acting as the "one voice" of the European HPC industry in relations with the European Commission and national authorities. The creation of this ETP fits perfectly with a European Commission's recommendation made in its recent communication¹ about HPC.

The timing of forming this ETP is right for two reasons:

- The coming Horizon 2020 Programme calls for a solid and energetic engagement in defining the HPC research content.
- We are at a crossroads: while new business models demand a new way of deploying – and designing – HPC systems, the technologies used today will not support the requirements of the next ten years. Thus, new approaches, technologies and methodologies are required.

We must "act now". Otherwise, the societal and economic benefits of a vibrant HPC ecosystem will not be realised.

The ETP strongly recommends **launching a research program with an aim to develop European technology in all segments of the HPC solution value chain.**

¹ Communication COM(2012)45 final "High-Performance Computing : Europe's place in a Global Race" : "The EU Industry engaged in supply of HPC systems and services should coordinate research agendas through the technology platform and thereby create critical mass of industrial R&D in HPC."

The ETP4HPC is ready to work with other stakeholders to implement such a research program, in order to develop a European HPC technology ecosystem and to help European HPC achieve recognized leadership.

This can be achieved through a research program that builds on the current strengths, takes advantage of technological disruptions with a tangible and sustainable market potential, chooses technologies that meet the needs of important applications, facilitates the creation of start-ups and the development of the SME sector and utilizes synergies with other Information Technology sectors.

We recommend prioritizing research in the following break-through technologies:

- Architecture evolution
- Energy driven HPC
- Extreme parallelism
- Programming paradigms, algorithms, HPC libraries and compilers
- Resiliency
- System Software
- Scaling I/O and storage with processing
- New storage solutions for big data
- Evolution of HPC workloads

We believe that such a research portfolio will help Europe build a competitive and sustainable HPC industry and ecosystem, produce world-class solutions for both academic and industrial purposes, reduce its reliance on technologies from other regions, increase the competitiveness and innovation capabilities of other existing industries, and address grand societal challenges.

CONTENTS

Executive summary	2
contents.....	4
1. Introduction.....	5
1.1 The European Technology Platform for High Performance Computing: ETP4HPC.....	5
1.2 Use of HPC in Europe	5
1.2.1 Prototyping and simulation	6
1.2.2 Large-volume data modelling and analytics.....	7
1.3 Why does Europe need to act now?	8
2. The ETP Vision.....	9
3. The ETP Mission.....	12
4. Strategy and Principles	14
4.1 Strategy	14
4.2 Principles of operation	17
5. The current position of European HPC industry	18
5.1 HPC Market	18
5.2 Synergies with other market segments	19
5.3 SWOT analysis	21
5.4 Conclusions	22
6. Technological objectives.....	22
6.1 Technological domains in the scope of the ETP.....	22
6.2 Developing break-through technologies.....	23
6.2.1 Architecture evolution.....	24
6.2.2 Energy driven HPC	24
6.2.3 Extreme parallelism	25
6.2.4 Resiliency	25
6.2.5 System Software	26
6.2.6 Scaling I/O and storage with processing.....	26
6.2.7 New storage solutions for big data.....	27
6.2.8 Evolution of HPC workloads	27
6.3 Mid-term 2016 objectives and milestones	28
6.4 Long-term 2020 objective and milestones.....	28
7. Links with other initiatives.....	29
8. Conclusions and Recommendations.....	30

1. INTRODUCTION

1.1 The European Technology Platform for High Performance Computing: ETP4HPC

Key European players² in the field of HPC research have decided to form a European Technology Platform³ to define Europe's research priorities to develop European technology in all the segments of the HPC solution value chain. HPC will strengthen European competitiveness in almost all industrial domains and provide a key capability for future research and innovation.

The ETP is an industry led initiative that will propose a Strategic Research Agenda taking advantage of European industry strengths to increase the value created in Europe from future HPC systems and solutions. Currently the design of supercomputer solutions faces significant challenges such as management of the extreme parallelism experienced in HPC architectures and the reduction of power consumption. Addressing these challenges presents significant opportunities for European players to improve their position in the worldwide market.

To achieve these objectives the current consortium will set up an organization that will be open to any businesses, groups or individuals who have R&D activities in any aspect of HPC located in Europe. The goal is to bring together all the research forces in Europe including R&D activities of SMEs, European corporations, international corporations and research centres to benefit from their competences and to foster these capabilities by proposing an ambitious research plan to the European Commission.

This document is the Vision Paper of the ETP4HPC presenting the objective, the mission and the strategy of the members.

1.2 Use of HPC in Europe

HPC is a key resource for innovation and leadership of industry in Europe. HPC has boosted scientific research and development across a broad set of disciplines from physics, chemistry, and bioscience to engineering. Numerical simulation, the third pillar of scientific discovery, on par with theory and experiment, has become the key to competitiveness in virtually all domains of industry and services. High Performance Computing also allows us to address the most crucial problems for our future: energy, weather forecasting, climate change, water... Increasingly, HPC has deep societal implications, in particular in drug design, medical research or in the entertainment industry.

HPC can be used for the prototyping of systems, simulating complex phenomena and for the analysis of large amount of data made available by sensors or produced by human activities.

² The founders are the industrial companies Allinea, ARM, Bull, Caps Entreprise, Eurotech, IBM, Intel, ParTec, STMicroelectronics and Xyratex associated with the HPC research centres BSC, CEA, CINECA, Fraunhofer, Forschungszentrum Juelich and LRZ

³ <http://cordis.europa.eu/technology-platforms/>

1.2.1 Prototyping and simulation

HPC is widely used in **Design and Manufacturing**. Manufacturing⁴ is still the driving force of the European economy, contributing over € 6,500 billions in GDP and providing more than 30 million jobs. HPC has revolutionised the design and manufacturing of a wide range of products. It can not only help to produce better products, but can also dramatically reduce the time and cost of design and production. These benefits have been achieved by HPC simulations of the finished product with a reduced need for building prototypes and the ability to perform more tests on the numerical models than on actual products.

HPC therefore enables high-value products to be made and brought more rapidly to market whilst avoiding prohibitive developments costs. It allows better product definition and reduced risk though better up-front analysis and decision-making.

As an example, HPC is widely used in the automotive industry to design efficient and safe vehicles. HPC is the cornerstone for improving the protection for drivers and passengers with crash simulations. Aerodynamics simulations lead to improved shape reducing drag and CO₂ emissions. Combustion models allow the design of more efficient engines. End-to-end optimizations can reduce the weight of the vehicles while maintaining their safety and comfort.

The European aerospace industry is also an important user of HPC. Cost reductions obtained in the development phase are even more spectacular compared to other sectors. The use of simulation tools has improved aircraft fuel efficiency dramatically; the fuel used per passenger mile today is around 30% of what it was 40 years ago. Simulations are also widely used for reduction of aircraft noise and to optimize airplane traffic.

Besides manufacturing industry, HPC is now also a key technology for the pharmaceutical industry. To design new drugs, the screening of molecules with specific features and shapes is done with the help of HPC systems. Before testing new drugs on animals or humans, HPC applications search for side effects using models of molecule and protein interaction.

HPC is also increasingly used in **Services**. Service providers need to simulate models or to perform optimization over complex problems and so require large computing power to deliver efficient services.

For example, in the financial and insurance sectors, HPC is used for estimation of the value of assets, pricing of products and risk assessment.

In the digital media industry, HPC systems are used to produce computer generated images. The availability of large computer power allows the development of new types of movies that can be used for entertainment, communication or training.

⁴ Analysis coming from PlanetHPC “A Strategy for Research and Innovation through High Performance Computing”

The energy sector is also an important user of HPC. For fossil energy, the discovery of gas and oil fields and the optimization of known field exploitation are the primary usages of HPC. HPC is also a key tool for the research of new energies like nuclear fusion.

1.2.2 Large-volume data modelling and analytics

Besides simulation of systems or phenomena, HPC is also becoming the main means for the analysis of large-volume data. The development of sensors, of measurement systems and of human data production will lead to an overwhelming volume of data that can only be modelled and analyzed with the help of HPC.

The best example is the genomic field where the detailed genomes of millions of species or human beings will soon be available. Only large computing resources can help to analyze these data and try to make them useful for understanding biological mechanisms. The HPC applications in this sector can lead to a better risk assessment of diseases, to customized therapies and to more efficient drugs.

The optimization of energy grids is also a challenge that can only be solved with the help of HPC. The availability of data about the user and system consumption in a large electricity network can lead to a gigantic optimization problem. HPC systems, able to quickly analyze the data and to propose an efficient solution for electricity production and distribution, will save energy and reduce the carbon footprint of networks.

Complex transportation systems (air or ground) can also be optimized through the use of HPC. The availability of sensors to track vehicles and of communication to influence their behaviour enables a control system that can optimize time to destination and fuel consumption. However, such a control system will need huge computing power and data bandwidth.

HPC can also be used for environmental risk prevention or management. Better weather forecasting will help to anticipate and plan efficient actions in agriculture, tourism and event planning. Global warming is also an important area in which HPC can have an impact. More accurate climate modelling and simulation will improve the understanding of natural phenomena and help to plan actions that can have a real effect in the long run.

The use of HPC for more energy efficient systems contributes to reducing the carbon footprint of human activities and helps address the global warming problem. The management of environmental crises (earthquake, tsunami, flood, fire) can benefit from HPC systems. Fast simulations taking into account the first data from a catastrophe can help to understand its likely evolution and to plan emergency responses to minimize the effects of the crisis.

1.3 Why does Europe need to act now?

The examples given above of existing or potential uses of HPC show that Europe needs HPC for both industrial competitiveness and to address societal challenges. HPC is without doubt a key enabling technology for Europe's competitiveness. As mentioned in a recent report⁵ written by IDC for the European Commission "Today, to out compute is to out-compete".

Many countries worldwide are investing in HPC and some like the USA, China and Japan are investing large sums of public money in related infrastructure and technology. Europe needs to be at the front line of HPC to retain a competitive research community and an industry able to innovate and to produce new knowledge and products.

The whole economy of the European Union will benefit from a strong European HPC ecosystem that will support the more advanced use of HPC. An ambitious HPC development strategy in Europe will be translated into a competitive advantage over other world regions. Such HPC strategy is not only motivated by industrial concerns. The development of HPC will also have big impacts on several challenges among the most important for European society.

Although industry has come to rely on a continual increase in computing performance, the future of high performance computing will not simply be a continuation of the past two decades. Chip design intellectual properties (IPs), components, architectures and programming models will need to undergo revolutionary changes.

At the same time, new business models will arise, and HPC systems will extend their reach to new fields like business analytics in the enterprise computing area.

As stated also in multiple sources^{6,7} we stand at a crossroad in HPC. Action is required to keep the societal and economic benefits of HPC from stalling in the foreseeable future. Investments on multiple fronts are necessary, of which the generation of a European HPC technology ecosystem is one very essential one.

Also, Europe needs to remain an attractive region for research as this activity is essential to prepare for the future. The existence of a strong HPC ecosystem will be a valuable asset to motivate companies and managers to establish their research in Europe. More globally the educational system needs to interact with the HPC sector to develop a virtuous circle where more HPC trained people means a stronger HPC field and a stronger HPC field means better educated people to tackle tomorrow's important economical and societal challenges.

As an industrial forum, the ETP believes that this is the right time to invest in HPC technology development with a high return on investment generated both by the growth of the added value created in Europe and by increased industrial competitiveness thanks to the use of HPC. The involvement of industry will facilitate the development of the innovation which is a priority of the Horizon 2020 program.

⁵ IDC report "A strategic agenda for European leadership in supercomputing: HPC 2020"

⁶ PlanetHPC "A Strategy for Research and Innovation through High Performance Computing"

⁷ IDC Special Study, July 2011, # SR03S

2. THE ETP VISION

As explained in section 1, High Performance Computing has become a key element for the competitiveness of knowledge based economies like the European one. Industry needs significant computing power for the design of complex systems and to accelerate the introduction of innovative products and services to the market. At the same time, scientific research relies more and more on simulations to produce new knowledge. HPC can really make a difference in many ways like more efficient transport systems better exploitation of resources (oil industry), development of new resource. This technology is becoming an engine of the economy. Its development is essential and its dissemination key.

Europe needs to have the infrastructure to access various computing systems covering a large spectrum of needs. It is important both to have widespread access to centres with an average computing power and to have a policy for providing access to the systems with the highest performance which are the only ones able to achieve breakthroughs. Nevertheless, a policy which focuses only on installing HPC systems, will not give Europe a leading edge in HPC as this domain requires comprehensive expertise and collaboration between technology providers, computing centres and users to make the most out of this technology. Only regions with a strong HPC ecosystem will be in position to make a difference through the use of this technology.

The creation of a strong HPC technology ecosystem is therefore very important. Europe has an opportunity to strengthen its position in this field and to take advantage of new challenges and disruptions brought by the evolution of technology. This is the right time to invest in the existing strengths of the European actors, like energy efficient solutions, architecture and integration skills, storage, system software and tools for HPC systems and HPC software, in order to develop European technology in all the segments of the HPC solution value chain with an ambitious research plan.

The exascale computing challenges are a good target to drive an ambitious research plan. Nevertheless, for an optimal return on investment and a sustainable development, we must keep in mind that a strong ecosystem can only be built if the technologies developed can address a broader market than just the top supercomputers. The strategy must on one hand give priority to technologies useful for the total current HPC market and even extensions of it and on the other hand take advantage of the investments made by other IT domains.

If such a research plan is put in place, the result will be at first a European HPC industry that is competitive worldwide and a strong ecosystem (made up of research centres, SMEs and large corporations) able to cooperate to increase the added value created in Europe of HPC solutions. The second consequence will be a comprehensive and high level expertise that will help HPC users, both research and industry, to exploit the potential of HPC. Therefore Europe will be in the first place to take advantage of HPC technology to solve societal challenges and to improve the competitiveness of its economy.

Our Vision for the European HPC environment is expressed in the following objectives:

1. To build a European world-class HPC technology value chain that will be globally competitive.
2. To achieve a critical mass of convergent resources in order to increase the competitiveness of European HPC vendors and solutions.
3. To leverage the transformative power of HPC in order to boost European competitiveness in science and business.
4. To expand the HPC user base, especially SMEs (through facilitating access to HPC resources and technologies) and to open the possibilities for SMEs to participate in the provision of competitive HPC technology solutions.
5. To facilitate the provision of innovative solutions to tackle grand societal challenges in Europe such as climate change, better healthcare, predicting and managing large scale catastrophes and energy-efficiency.
6. To foster international cooperation in research and industry.

1

Only a European world-class HPC technology ecosystem covering the complete value chain will enable Europe to be at the leading edge in HPC. In fact, a strong expertise in HPC technology is essential to developing the best practises in HPC and to anticipating the future evolution of this domain. With the HPC and IT technology evolution, this expertise will be more and more valuable for successful exploitation of HPC systems because of the growing complexity of this domain. Each link in the value chain is important as only a complete understanding of the available technology options can lead to choosing the optimal path for the development of HPC. It is also critical to build a chain that is globally competitive. So the strategy will be to invest in European strengths and to develop alliances wherever Europe has a lack of leadership with a goal to bring the needed skill or expertise to Europe.

2

The development of this European HPC technology ecosystem will mean more research and development resources based in Europe covering the complete HPC value chain. The objective is not only to have the expertise and skill to make the right choices and to develop the best HPC practises but also to exceed a critical mass of resources that will improve the global competitiveness of HPC solutions integrating technologies developed in Europe in the worldwide market. The effort of increasing the European technology ecosystem must have a positive economic impact for Europe. The added value created in Europe must increase in the HPC market and generate benefits in Europe to sustain the development of this ecosystem.

3

With the existence of a strong HPC technology ecosystem, Europe will be in the best position to take advantage of HPC to increase the competitiveness of its science and business. The economic success of European HPC technology players will per se be an achievement. But it will also have an impact by increasing the potential for users both in research and in industry to cooperate with expert teams of technology providers and to improve their usage of HPC. The presence of a vibrant HPC technology ecosystem will increase the general awareness of HPC and help its diffusion into new sectors. The existence of expertise and skills covering the complete value chain from technology components to usage will be beneficial to all actors in the HPC field.

4

HPC is already widely used in research and in some industrial sectors. Nevertheless, other sectors could also benefit from the usage of HPC, especially SMEs. For SMEs, availability of HPC resources is important for their competitiveness and for better adaptation to market demands especially in terms of innovation and fast renewal of their product and service offerings. The development of new HPC delivery solutions such as cloud computing, of new HPC services helping newcomers to master these technologies and of easy to use HPC solutions is a priority to expand the usage of HPC and to spread it across industry including SMEs.

5

While building the HPC technology ecosystem, one objective will be to tackle grand societal challenges that Europe is facing. Priority will be given to actions that could have an impact for European society. As explained in the previous section, HPC has the capacity to contribute to a new era in healthcare systems, to produce important information to make decisions for sustainable development or to develop more energy efficient systems. The ETP4HPC will promote actions that will go in these directions. Public support for the development of the HPC technology ecosystem will require a return on investment e.g. through addressing critical challenges Europe is facing.

6

The development of the European HPC technology ecosystem will be an opportunity to increase international cooperation. Many countries or regions have decided to invest in HPC and we must avoid “reinventing the wheel”. The effort in Europe must be open to cooperation but needs to be strong enough to have discussions on an equal footing. Only mutual cooperation will be beneficial to all the actors. This cooperation will of course cover the research agenda but could also lead to industrial cooperation that could help the worldwide diffusion of technologies developed in Europe with an appropriate return on investment for Europe and a fair share of the IP value.

3. THE ETP MISSION

The mission of the European Technology Platform in the area of High Performance Computing is to unite the efforts of various European industrial, research and end-user organizations with the aim of producing world-class HPC hardware and software technological solutions along the entire HPC value chain through:

- Designing and updating a Strategic Research Agenda (SRA) to provide decision makers with relevant advice and expertise for the long term development of HPC in Europe.
- Providing recommendations and support to the implementation of the SRA.
- Facilitating coordination between the HPC ecosystem and public authorities (EU and Member States) responsible for HPC research and dissemination programs.
- Fostering joint initiatives among ETP members and other stakeholders in the area of research and innovation programs.
- Facilitating the emergence of start-ups and the growth of existing SMEs.
- Supporting Europe and Member States authorities by reinforcing Europe's position in the worldwide HPC arena.
- Representing the voice of the European HPC industry in the worldwide HPC arena.

ETP4HPC will design a Strategic Research Agenda (SRA) and will update it on a regular basis to take into account the ongoing evolution of technology. This document will provide an analysis of the technological context of HPC and propose an action plan to reach the objectives set up by the ETP and presented in the previous section. It will be inspired by the strategy (developed in the next section) written with the participation of all the organizations having HPC research activity based in Europe and willing to contribute. The recommended action plan will target long term objectives needing some continuity of commitment but also some shorter term objectives associated with specific milestones. The SRA will be presented to the European Commission with the objective of influencing future calls for projects in line with the proposed action plan. The European Commission will manage the calls and the ETP will be open to providing advice and expertise when needed.

As an industrial forum, the ETP will provide recommendations on the implementation of the SRA. These recommendations will include of course research priorities but could also deal with educational issues or legal framework proposals (e.G. tax measures) that could enhance the development of HPC in Europe. The targets of these recommendations will be European organizations but also any regional or member states organizations that could have a positive impact on the development of HPC. In addition to these recommendations, whenever it will be possible and according to its means, the ETP will support the actions in line with its SRA. This could mean participation in actions for the development of HPC awareness or for HPC dissemination.

The ETP will gather a strong community of HPC stakeholders ranging from SMEs, corporations, research centers to universities. While preserving the complete freedom of its members, the ETP will coordinate its member activities whenever necessary to facilitate the relationship with public authorities. This could be providing inputs on a question of interest to the EU or some Member States, for supplying relevant information about the HPC domain or for organizing the participation of the HPC ecosystem in dissemination programs. The ETP will also provide a communication channel to the HPC industrial and research community.

The ETP will also serve its community by helping to create joint initiatives among its members in the area of research and innovation programs. The ETP, of course, will not interfere with the creation of consortia but can facilitate the emergence of complementary teams (e.g. of large and small companies, research centres and users) by organizing brokerage events to prepare research project proposals. By gathering different kinds of HPC actors, the ETP will increase the interaction of the European HPC ecosystem and make it more active in finding new cooperation frameworks that can be beneficial to its participants. The active encouragement of the ecosystem will be an important task for the ETP operational team that will be put in place to carry out these missions.

The SMEs will play a critical role in this ecosystem. Creation of start-ups is an efficient means to transform research or innovation idea into market values. The entrepreneurs (coming from the academic world or the industry) need to find a favorable environment to create companies and to develop them. Existing SMEs need also to be able to develop themselves and benefit from an environment with adapted services, good relationship with academic research and support of large companies to facilitate their growth. In Europe, HPC software ISVs are often SMEs. As they have limited resources to adapt their applications to the technology evolution it is important to provide them with means to overcome this issue and to continue their growth. The ETP will facilitate the creation of this favorable environment for new and existing SMEs.

The existence of the ETP will be an asset for the European Commission and the Member States in the worldwide HPC arena. At a time when a lot of countries are investing in HPC, the presence of an industrial forum can only reinforce Europe's standing showing that our region has strengths and that the players are ready to coordinate their efforts in building a stronger HPC ecosystem. The ETP will show the domains of European excellence in HPC and the efforts being taken to strengthen or to expand them or to supplement them to fill any gaps. It will facilitate cooperation on an equal basis between Europe and other countries or regions.

One mission of the ETP is also to become a recognized organization that can speak for the benefit of the European HPC technology industry. In the global context, it is important for the players having HPC research in Europe to be able to express their point of view on many topics. The ETP will federate the voices of its members and give to them the ability to be heard more strongly and clearly. For example, if some cooperation at international level is set up on the subject of exascale, the existence of the ETP will be a chance for the European actors to get more recognition and influence in defining the orientations and the roles of each region. The ETP will work to build international recognition of the potential of Europe as a provider of HPC technologies.

To develop all the facets of its mission, the ETP will set up an operational team that will act on behalf of the ETP members.

4. STRATEGY AND PRINCIPLES

4.1 Strategy

The Vision and the Mission will be implemented according to the following strategy:

- The ETP will collaborate closely with scientific and industrial HPC user communities and ISVs in Europe to ensure the Strategic Research Agenda addresses the main HPC challenges.
- The Strategic Research Agenda will include cooperation with other existing initiatives in the area of HPC and ICT in Europe such as PRACE.
- The Strategic Research Agenda will specify the most effective actions to take advantage of Europe's current strengths, utilise existing and emerging opportunities and close the currently increasing gaps. Going beyond those initial priorities it will also identify strategic industrial and research user requirements to define the research themes for the next generation of key longer-term technology innovations.
- The ETP will take into account other key enabling technologies from non-HPC industry and research centres that might bring disruptive solutions into HPC.
- The ETP will facilitate the creation of new international standards if required in order to expedite the deployment of a globally competitive HPC ecosystem.
- The ETP will recruit new members such as vendors, ISVs, research organisations and end-users with active R&D in Europe.
- The Strategic Research Agenda will focus on the main challenges facing the European HPC industry such as: parallelisation of software, power consumption, reliability and big data.

The ETP will focus its action plan on the development of HPC technologies. Nevertheless, the technologies must be aligned with the needs of the applications. It would be meaningless to propose investments in technologies that will not at some time be useful to applications. This is why the ETP will closely collaborate with the user communities and the application ISVs (Independent Software Vendors) to be sure that the SRA priorities will be correctly aligned with the development of the future HPC applications. The user communities and ISVs will also be keenly interested in the actions proposed by the ETP on topics such as programming models and tools for application development. Some interactions around “co-design” activities will also be developed. The user communities will contribute with application requirements for more efficient computation and the technology providers will explain the technology trends that need to be taken into account while preparing new applications. So there will be close interaction between the ETP stakeholders and the user and application ISV communities.

The SRA will also take into account the synergy with other initiatives related to HPC. Section 7 will explain in more detail the vision of the ETP regarding interaction with other initiatives. The strategy will be to work in close coordination and in a complementary way and never with a competitive spirit. For example, the ETP sees a fruitful cooperation with PRACE. PRACE develops the HPC infrastructure for European researchers and works on education and training. ETP4HPC can complement this action by bringing expertise and skills for making the most of the infrastructure and by advising on the technology trends in infrastructure and application development.

In implementing the SRA, the strategy will be to take into account the strengths, the existing and emerging opportunities and the objective to reduce existing gaps. Only actions that can lead to a sustainable effort and to the emergence of technologies that will find a place in the market will be supported by the ETP. Priority must be given to actions that have a chance to produce competitive technologies in the worldwide arena. To cover the complete technology value chain, the options can be different depending on the situation of European players. On some subjects, the most effective way could be to work in cooperation as long as the exchanges can be fruitful for all the parties. In some other areas with existing strengths, Europe could reinforce its position to achieve worldwide leadership. In any case, the objective will be to increase the use of technologies developed in Europe within globally competitive HPC solutions, with a fair return on the added value created in Europe.

To design the SRA, the ETP will carefully look at the potential of the whole HPC market but also of other IT markets. The idea is that although the HPC market is an interesting one, it is still a small one compared to other IT markets. So, it makes sense to either develop for HPC a technology that can span other IT markets or to take advantage of technologies that will be developed by the other IT market forces and to use them in HPC without the need to finance their development. In both cases, the economic benefits for the HPC community will be important. Some past examples show that the HPC community can benefit from disruptive approaches coming from other IT domains. On the other side, the ETP will be very careful in defining the priorities in order to avoid investing in technologies that will be unable to reach sufficient market coverage to generate a return on investment. Even if this strategy leads to Europe reaching exaflops performance later than other countries, we can not afford to invest money in non profitable technologies.

One other feature of the strategy will be to facilitate the creation of new international standards. The emergence of a standard is often important for a new technology to reach market acceptance. The potential users of a new technology could be reluctant until they are reassured by the existence of a standard that means their investment will not be wasted. The ETP as a group of stakeholders has a role to play in promoting new standards. The ETP plans to work with other organizations to cooperate on the creation of new standards for HPC technologies. This activity will be important for example in the field of programming models where application developers most of the time can not afford to invest in solutions that will not be supported by a standard. The creation of international standard is also a way to develop a more competitive HPC ecosystem where the players can find a solid framework in which they can invest more securely.

As will be explained in the Principles of the ETP, the SRA will be designed by a large community of members that will come from technology providers, vendors, ISVs, research organization or end users with active R&D in Europe. The key European players that have taken the initiative of establishing the ETP and publishing this Vision Paper will be reinforced by other stakeholders to cover the whole spectrum of topics and of areas of interest. They will manage the ETP during the bootstrap phase and thereafter the ETP will operate according to its organizational rules. So the SRA will be the result of a wide consultation of HPC actors.

The SRA will focus on the main technical challenges facing the European HPC industry. Section 6 will develop in more detail some of the technical challenges but the main idea is to focus the SRA action plan on the most crucial topics for the development of future systems and applications. These will include the exploitation of extreme parallelism that will come with future HPC architectures, the optimization of power consumption of systems, solutions for increasing the reliability of large systems and the management, analysis and processing of very large data sets.

This strategy will enable the ETP to fulfill its objective and mission. It will lead to a SRA that will be relevant for research excellence, economic development, European HPC technology ecosystem competitiveness and development of HPC in Europe.

4.2 Principles of operation

The ETP will be open to any organization wishing to join and that has HPC research based in Europe. We expect to get members coming from the complete value chain of technologies involved in HPC systems (hardware, architecture, system software, tools, and application software). All Businesses or Institutions in these fields doing research in Europe will be welcome. HPC users who are active in HPC research will also be invited to join the ETP as members.

The organizations not involved in research but interested to participate in the ETP, would be able to apply for the status of Associate Member. If accepted by the ETP, the Associate Members will be involved in the discussions of the ETP and will be able to participate in some activities.

The ETP will operate in an open, transparent and democratic manner according to its organizational rules characterised by trust and accountability among all members. It will not restrict members' freedoms:

- Freedom of association: All parties involved may associate and collaborate with whomever they want, regardless of geographical origin, name, business size etc.
- Freedom of competition: The ETP will not comment on or otherwise influence any research proposal, commercial or pre-commercial tender. Members will compete as usual, and the ETP will have no role, publically or privately, in judging or commenting upon any particular competitive situation.
- Freedom of communication: Participation in the ETP in no way limits members from communicating their own positions regarding HPC, provided only that they do not misrepresent any formally agreed position of the ETP.
- Freedom of IP management: No member will be restricted in any way from developing or using IP for which they have rightful ownership nor will they be required to share or otherwise open such ownership.

The ETP will set up an operations team. This office will be representative of the European nature of the ETP and will operate as an international team. It will be the main resource of the ETP to fulfil its mission beside the commitment and the involvement of all the ETP's members.

5. THE CURRENT POSITION OF EUROPEAN HPC INDUSTRY

5.1 HPC Market

The value of the global HPC market, as defined by IDC in its regular reports, includes the total values of the purchases of HPC systems (ranging from technical computing servers up to supercomputers) as well as the value of associated storage, middleware, application software and services. According to a study made for the European Commission⁸ the worldwide market value in 2009 was €12.9 billion and the forecast for 2013 is €17.3 billion.

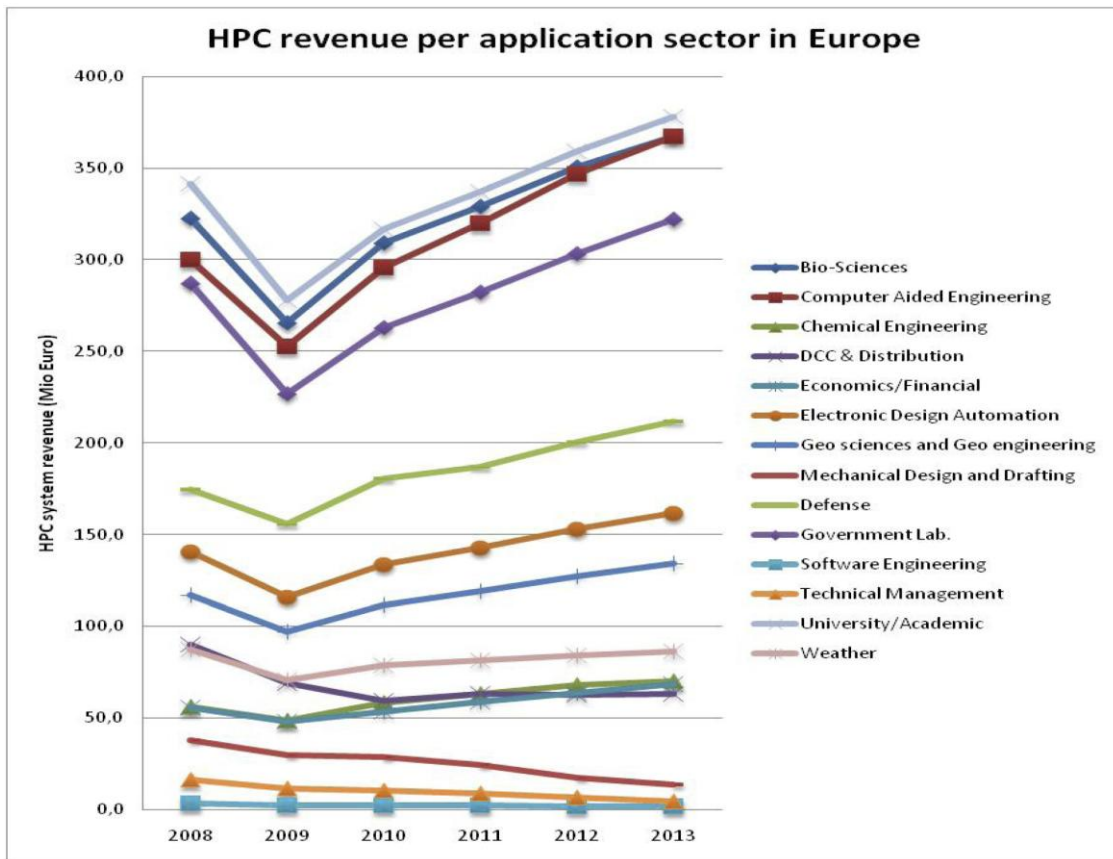
The US is the largest market (almost 50% for supercomputer) but other regions have important growth rate such as Asia or the BRIC countries (Brazil, Russia, India, China). This situation need to be taken into account to achieve the objective to increase the added value created in Europe in the HPC solutions sold worldwide.

The European market represents just over 25% of the worldwide market: its value for 2009 was €3.3 billion and the 2013 forecast is expected to be €4.6 billion.

The HPC system category is growing less than the other categories. One of the more dynamic segments is 'storage systems' due to the 'data explosion' associated with the growth of large and complex simulations.

A forecast of HPC expected investments in the different application sectors of HPC systems is shown below. The largest sector is 'University/Academic' followed by 'Biosciences', 'Computer-Aided Engineering (CAE)', 'Government' and 'Defence'. The highest growth in the coming years is expected to come from the 'Computer-Aided Engineering' sector.

⁸ 'A strategic agenda for European leadership in Supercomputing : HPC2020' September 2010, IDC #SR03S



European Union - HPC Revenue (€000) Forecast by Application, 2008-2013 (Source: IDC, 2010)

5.2 Synergies with other market segments

The HPC market is the most dynamic sector of Information Technology. A recent study predicts that the industry will continue to grow at a 7.0% compound annual rate. Above-average growth is predicted to continue in 2012 before moderating to sustainable rates later. We believe this is a ground-breaking trend. HPC is becoming a model applicable to all types of high-end computing. Also the differences between High Performance Technical Computing (HPTC) and High Performance Business Computing (HPBC) are becoming less obvious:

- Typical HPC technology is used today in logistics, entertainment industry, financial services, and insurance data-mining over structured and unstructured data, which can hardly be considered as scientific and technical workloads;
- Traditionally restricted to using moderate multi-threading parallelism, business applications increasingly utilise all kinds of parallel computing, including SIMD processing;
- On the other hand, many typical scientific workloads (e.g. DNA sequencing) do not require particularly high floating point performance;
- Interestingly, both HPTC and HPBC use the same processors (for instance the ubiquitous X86-64 architecture). Even though business applications make lighter use of the floating point units, it is simply more economical to rely on the same high volume, low cost components.

Thus, there is evidence that high-end Scientific/Technical Computing and high-end Business Computing are converging on multiple common technologies. Furthermore, the wave of downsizing initiated by the emergence of COTS-derived low-cost, high-volume servers is rapidly mutating in an enormous movement of re-concentration. There are several reasons for that:

- Modern microprocessors have become so powerful that assigning a server to a single service or application does not make economic sense. Nowadays, even the smallest microprocessors tend to be multi-core ones and will continue to evolve to act as multiprocessors or even clusters-on-a-chip. Not surprisingly, users have resorted to gathering many aging servers into a single unit;
- The virtues of physical isolation provided by downsizing were nevertheless preserved by using virtualization techniques. At the expense of a minimal overhead, hardware assistance and hypervisor-based system software ensure that the malfunction of one application or the demise of a service would not propagate to the whole system. Virtualization allows us to answer all needs with large servers and thus eliminate the “one size does not fit all” issue;
- The resulting pooling of resources is beneficial from an economic as well as environmental point of view. It brings benefits of scale to all logistical elements of the problem: power supply, cooling, floor space. Large savings are also to be expected from the manpower required to manage such aggregate implementations: optimized administration, monitoring, maintenance;
- Last but not least, the changes in the world’s economy also require evolution in the use of IT technology. Users want to concentrate on their professional tasks without being bothered with tasks that are not part of their core competence. The trend is not new. Partial attempts to meet these requirements were proposed through outsourcing, and later on with on-demand computing; however, they have only recently materialized as the ‘As a Service’ concept. Together with the emergence of ‘cloud computing’, there is an assumption that in the near future, users will access IT as a service.

All these considerations lead us to draw the following simplified picture of tomorrow’s IT:

Our universe is populated with billions of smart devices, increasingly mobile, increasingly connected using the Web standards. Often referred to as ‘Internet of things’ this multitude of terminals (smart phone, tablet, PC), sensors, RFID tagged objects, industrial or household appliances rely on Internet protocols for access to remote computing and storing data, making computing an inseparable part of our everyday activities while simultaneously disappearing into the background.

On the other hand, as a result of concentration, as explained above, large data centres and HPC computing centres offer huge computing power and storage capacity for cloud computing services or heavy simulations. HPC technology has some synergy with this IT segment that goes well beyond the domain of applicability of the classical sense of HPC.

The ETP will exploit this synergy to increase the sustainability of the development of HPC technologies. Whenever possible the technologies developed for HPC will address this larger market. On the other hand, technologies appropriate for HPC coming from this market will be incorporated to minimize R&D investment.

5.3 SWOT analysis⁹

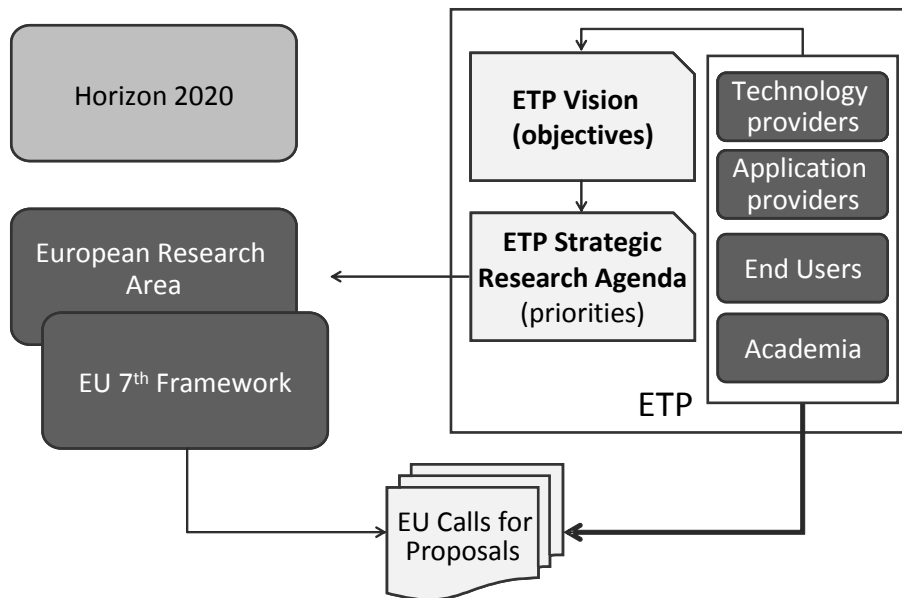
POSITIVE IMPACT	NEGATIVE IMPACT
<p style="text-align: center;">Strengths</p> <ul style="list-style-type: none"> • World-class expertise in technologies that are crucial for future HPC systems: e.g. low-energy devices, software tools. • Some established industrial players providing world-class HPC systems and solutions, process architecture, software tools, HPC allocation software and high-end storage solutions • A world-class scientific and engineering community • Established community of research centres • Established processes of knowledge transfer to industry • Expertise in certain advanced technology areas (e.g. embedded systems, semi-conductors, communication technology) • Vibrant Small and Medium Business (SMB) sector • Success of a number of cross-regional industrial programmes (e.g. Airbus) • Established educational programmes and expert communities 	<p style="text-align: center;">Weaknesses</p> <ul style="list-style-type: none"> • The market share of European HPC solutions is relatively small on both global and European level • Fragmentation of the European resources • Limited internal market size might hinder HPC growth unless the industry is competitive globally • Under-investment in HPC (while other regions are ramping up) • Many European HPC companies are small and have a limited investment capability (in particular HPC ISVs) • Many European companies consider HPC as a tool only without looking to use it in order to gain competitive advantage • European scientists consider European research programmes uncompetitive in comparison with those of other regions • Lack of next generation HPC scientists • Shortage of qualified HPC job candidates • Financial crisis in some of Europe's economies might affect the availability of research and development funds
<p style="text-align: center;">Opportunities</p> <ul style="list-style-type: none"> • Technical trends heading in directions where Europe is strong: e.g.: explosion of parallelism that requires new programming models and software tools, reduction of power consumption leading to the use of low power devices, data explosion means efficient storage and parallel file systems; HPC expansion to some time-critical applications required by embedded systems. • HPC solutions have the potential to strengthen Europe's most competitive industries • The shift to exa-scale HPC computing will require new technologies • Software applications will need to be rewritten in order to meet HPC requirements • Deployment of significant funds within the EU aimed at developing HPC infrastructure and technologies • Synergies with other IT market can be exploited 	<p style="text-align: center;">Threats</p> <ul style="list-style-type: none"> • Other regions are investing heavily to develop HPC technologies and could overtake Europe and assume a leading position • Other regions could invest more in HPC infrastructure than Europe and so out-compute and out-compete both in research and industry • The development of aggressive HPC strategy in different regions can lead to a brain drain leaving Europe without critical HPC skills • While the European market is completely open, various form of protectionism in some regions can reduce the chance of European HPC technology providers to increase their market share

⁹ See also an interesting analysis of HPC player position made by the EESI project: http://www.eesi-project.eu/media/download_gallery/EESI_D4.3_WG4.1-report_R1.0.pdf

5.4 Conclusions

European technology providers, applications providers, end-users and research centres should collaborate in industrial research projects with an aim to achieve global leadership in the project domains.

This could be achieved through the work of a European Technology Platform:



6. TECHNOLOGICAL OBJECTIVES

6.1 Technological domains in the scope of the ETP

The objective of the ETP is to develop European technology in all the segments of the HPC solution value chain. Thus, the scope of the ETP will include hardware, system software, application software and HPC related services. These domains will be addressed with a global technological approach as the ETP will propose to develop technologies that will form a basis for building solutions in these fields.

The following topics will be in the scope of the ETP:

- Microprocessor, accelerator, memory architecture, system and SoC (System on a Chip) architecture, interconnects, storage controllers, storage architecture...
- Operating system, system management, run-time system, scheduler, resource management, cluster management, file system, check-pointing methods...
- Programming model, compiler, debugger, performance and energy monitoring tools...
- Numerical libraries, application programmability, data visualization, scientific data management...
- Application optimization, exploitation optimization, computing centre energy management, problem modelling methods...

Application development itself is outside of the scope as it primarily lies in the domain of the applications and covers all disciplines of science and technology (it is not HPC technology per-se). Within the scope, however, there is the provision of enabling technology to make such application development and deployment as effective and efficient as possible.

The ETP will closely work with the application communities and the application ISVs. We foresee co-design projects looking at the links between technology capacities and application requirements (computing patterns, memory bandwidth, interconnect bandwidth, communication pattern, IO pattern, IO bandwidth...). Another field of cooperation will be around the assessment of the results of the ETP's program. The ETP should make sure that it targets relevant applications. It might be interesting to choose some applications important for Europe's challenges and to use them for the evaluation of new technologies. This could mean having pilot projects aiming to illustrate the usefulness of the developed technologies.

6.2 Developing break-through technologies

The ETP will focus mainly on technologies that can bring answers to key challenges. Computing performance is now driven by multi-core processors and tight integration of system components. An increase in computer performance will be achieved through an increase in core count, in core performance and closer integration but also through software tools capable of taking advantage of the large number of cores. Europe can take advantage of this new paradigm, building on the strengths of European players. The exascale challenges are appropriate targets to set ambitious goals as long as the needed technologies can also be useful for a large part of the HPC market and even better be used in other IT domains.

We can expect that most technologies relevant for exascale computing would also be relevant for solving less complex problems. Nevertheless, they must be cost effective when applied to less powerful solutions and avoid high portability costs to run applications at a lower scale. The ETP will be careful in proposing research directions that fulfil these conditions.

On the technological landscape, major changes are expected. Traditionally, microprocessor frequency increase has driven the ever increasing computer performance. From 2006, additional performance has been gained by adding more processor cores, at the same or lower frequencies. The first wave of technologies exploiting these multi-core environments is now fully optimized. Getting to the next level of computing power (x1000) requires a whole new computing environment, supported by breakthrough technologies both on the hardware and software side. This is an open field and there is today a clear window of opportunity for Europe to become an active player in these new technologies through proactive R&D programs.

The next paragraphs present some of the challenges that will be priorities for the ETP's research plan.

6.2.1 Architecture evolution

Server architecture will change because of technological opportunities and new requirements coming from the system level.

The miniaturization of CMOS will continue providing a larger transistor budget. Chip designers will have several options for using this budget to increase the overall performance. The search for the right balance between increasing the core count, providing optimal memory bandwidth and improving communication could lead to new architectures. Three-dimensional integration will also be a technological disruption that will influence the architecture giving more opportunities for innovative processor-processor and processor-memory integration. The integration of optical and emerging resistive memory technologies will also be a big change that can allow new solutions at the server and even the HPC system level.

The architecture designers will certainly face challenges that will influence future systems. Power efficiency (see below) will be a key driver for the selection of the right solutions. Low power technologies need to be adapted to HPC because local power optimizations are not enough to get global system power efficiency. The trend to increase the parallelism of the architecture represents another challenge for the architects: how to provide environments not too difficult to use at the system level. This question will be at the heart of architecture research in the next years. Big data applications might also mean changes in architecture to provide new trade-offs between processing, memory and communication.

6.2.2 Energy driven HPC

Proposing new approaches to design energy efficient HPC solutions is a big challenge. Limiting power consumption and related thermal emission has become the leading constraint in designing large systems. Power consumption is a decisive parameter in calculating the cost of ownership of a system and large computing centres indicate that the power envelope for large systems should not exceed 20 MW.

Clearly power efficient processors and SoC are one key element in this challenge, but a system level perspective is also required. Adding these functionalities to systems need to be done without negatively impacting the overall system performance, power consumption, latency and avoiding creating additional jitter. This is a major challenge, to be solved both by hardware and software solutions.

The power consumption demanded at the system level by the application software is also a critical element. Re-architecting the software to limit its power consumption, developing new data management schemes to improve data localization are challenges of interest to the ETP.

6.2.3 Extreme parallelism

The increased parallelism of future architectures will require major improvements in system software and programming models while still providing a migration path from existing environments. Data placement in the memory hierarchy will become more and more important and may require explicit APIs to help programmers in their research of performance and scalability. New smarter runtime systems that can take decisions on behalf of the programmer may also be very useful. Similarly, optimizing I/O performance will barely keep surviving with legacy interfaces with file and storage management. More dynamic and application-driven interfaces will prove necessary to perform proactive caching, data reorganization and migration, real-time analysis and pre-processing.

Heterogeneity, growing in-core/inter-core performance discrepancies and the use of accelerators also raise the question of using a mix of several programming models; the temptation to unify on message passing is synonym of downgrading the capability offered by multi-core chips with shared memory. While existing applications could still run on such massively parallel systems, most demanding ones will no doubt need rewriting to take full advantage of the latest generation of computer potential. Compilers usually generate code adapted to individual nodes and need additional tools or runtimes to adapt to parallel programming models. With the advent of massively parallel heterogeneous systems, compilers must encompass additional awareness or rely on specialized tools. Another area requiring major efforts is the tooling for debugging and performance optimization. These utilities must scale at levels never envisioned before. One way would be to take advantage of the lesson learned by performance tools community by the means of development of APIs collecting states and values and using, among others, methodologies from signal processing and data mining.

6.2.4 Resiliency

Moore's law brings huge benefits by continuously increasing the transistor budget chip designers can rely upon. At the same time, miniaturization also is source of new issues. More transistors (and smaller ones), lower voltages to cope with energy saving requirements mean greater probability of failures. In parallel, the number of chips also grows, and interconnection networks become more complex and larger, spectacularly increasing bandwidth needs and lessening latency, resulting in very small noise margins. It is safe to predict that the mean time between failures of a large system will measure in minutes. Without dealing with this issue from the ground up in both the hardware and software architecture, resiliency will become problem n°1 for the HPC world, where weeks or months long runs are usual.

The first absolute requirement to correctly apprehend the problem is the implementation of adequate mechanisms for detecting and reporting hardware errors, such as system bus errors, ECC errors, parity errors, cache errors, translation look aside buffer errors (Machine Check Architecture in Intel's terminology), core failures. Data paths will need to provide alternate routing whenever a link becomes corrupted. System software will have to detect error conditions and provide redundant resource management and reconfiguration support to resume execution. Middleware will have to be redesigned with fault tolerance in mind

(e.g. fault-tolerant MPI offering a paradigm for resilient programming). Performance and tuning tools will need to cooperate with system fault tolerant interface features too.

It is clear that current state of the art based on global checkpoint / restart points of recovery are too coarse grained to keep on being the solution. A more local approach to rollback and recovery must be implemented, with standard APIs exposed to applications. Here, a significant interaction with application developers will be required.

6.2.5 System Software

System software will emerge from an ecosystem capable of developing a software stack in a co-design spirit with the hardware and application community.

The vertical integration of complementary system software components is a key challenge that will require proposing perennial public API (e.g. thread API, tracing API, ...) that also contributes to creating and enhancing international standards.

New programming models (e.g. PGAS) and high level abstractions will have to be provided to master the complexity of heterogeneous hardware while taking into account fault tolerance as well as energy efficiency. Many new compiler techniques, runtime libraries and performance profiling tools are to be developed to help code tuning and facilitate runtime adaptation to varying hardware resources (e.g. auto-tuning techniques). Finding the right trade-off between portability and high performance will be the main challenge.

New tools, software engineering techniques, methodologies are required to help legacy codes transition to extreme parallelism and their interactions with new codes and new algorithms.

The path to the design of a software stack relies on the capability of testing and evaluating new software and techniques. The availability of benchmark suites that exposes the main application characteristics is on the critical path of the software stack design and will ensure focused and efficient efforts of the entire ecosystem.

6.2.6 Scaling I/O and storage with processing

As processing power increases and as efficiencies in programming methods for extreme parallelism are developed, the performance of data storage systems required to support new computer systems must scale at even faster rates than the basic increase in FLOPS. This applies in particular to sustainable data rates and file creation rates. The primary reliance upon disk drive technology for non-volatile storage, with a well established roadmap typically providing around a 10% annual improvement in data rates shows that this cannot be relied upon to deliver the 100-1000x performance leap that is required.

Analysis of the whole Input / Output stack as deployed in typical HPC systems today shows that almost all elements of the system will need to be addressed. Significant performance is lost in the middleware layers, and it is critical to address the overheads imposed by the Posix based access model. Non-volatile storage media (solid state) must be seamlessly integrated into data storage systems if the required data rates and file creation rates are to be achieved. These changes are needed at all layers and the systems need to embrace new and unproven technologies (eg. Byte addressable devices such as PCM, Memristors etc). This can significantly

change existing programming models. Overall scalability of the file systems must be dramatically improved and efficient methods for migration and pre-emptive management of data must be developed. The raw capacity growth of disk storage is better matched to the needs of the future, though this has a number of technological hurdles to overcome if the growth is not to drop well below its 40% CAGR seen over the past decades (eg usage of technologies such as Bit patterned Media, and Heat Assisted Recording could be potential solutions). Overall, the total devices required within HPC installations will increase significantly making the challenges of power efficiency of such systems increase in significance as the storage begins to take a larger percentage of the total data centre budget. With such huge systems comes the inevitability of failures, as described in 6.2.4. Resiliency to failures without significant loss of performance is a critical requirement and must be designed into the basic architectures used within storage systems.

All of these areas are within the scope of technology developers and researchers within the European community.

6.2.7 New storage solutions for big data

We are continually bombarded with new opportunities to advance our science, knowledge and social capabilities through the analytical use of data. The widening range of digital data sources is well documented to be growing at exponential rates. Increasingly, traditional HPC is being complemented by new 'Big Data' processing to mine for key insights. This crossover between 'HPC' and 'Cloud' is expected to continue and it is critical that storage solutions are developed that reflect the very different nature of the storage access patterns to these huge datasets, be they structured, semi-structured or un-structured. Other characteristics such as data security and Quality of Service will be much more critical in such environments. These areas are not well covered by today's HPC storage systems. It is also critical to ensure a co-design of storage and computing methods / algorithms as demonstrated by the success of MapReduce (Hadoop) to enhance parallelism in the essentially serial world of data storage.

6.2.8 Evolution of HPC workloads

Most HPC workloads are actually composed of several separate processing and evaluation steps that are coupled together. Traditionally these steps are thought of as independent pieces of work that are 'coupled together' by the team that is doing the overall evaluation. The steps might be as simple as: 1. Input data capture or creation. 2. HPC program execution on the input data 3. Visualization of the results. However, quite often the workflow is much more complex and involves several execution steps at each of the phases.

In the upcoming era of very large data sets being created at each step it is no longer practical to think of each of the steps as loosely coupled, independently scheduled events. APIs and tools will need to be created to allow the users to describe the entire workflow so that the global resource managers and job schedulers can be looking ahead to start moving data to the next piece or computation, or moving the next piece of computation to the data. It will not be practical to merely 'write the data out to file' and then 'bring it back in' for the next job step, and the life cycle of the data and execution state will have to be carefully managed.

This new era implies not just new APIs for describing the workflows, but substantially changes what 'resource management' and 'storage hierarchy/data life cycle management' means.

6.3 Mid-term 2016 objectives and milestones

The complete analysis, a roadmap proposition and the final selection of the milestones will be done during the design of the Strategic Research Agenda. The following could be interesting mid-term objectives:

- Low power computing technologies and a roadmap for future evolution.
- Energy efficient computing technology roadmap.
- Highly scalable IO and interconnect/networking system.
- New layers in the HPC stack mixing OS and middleware functionalities focused on energy efficiency, performance optimization, improved resilience, etc
- Software ecosystem (compiler, library) related to energy efficient computing technologies, improved application design productivity, ease of use.
- A suitable programming model for heterogeneous and massively parallel architectures.
- API for memory hierarchy and thread management, power management and inter-node communication.
- Software API for reliable system and application development.
- Scalable performance analysis and prediction tools for exascale platforms.
- Management system able to deal with energy optimization.

The APIs will allow development of both hardware implementing the functionalities and software that can use these mechanisms in a co-design spirit.

6.4 Long-term 2020 objective and milestones

While the SRA will discuss these in more detail, the following objectives could be proposed for 2020:

- Demonstration of an exascale system based on technologies developed by the ETP ecosystem.
- European solutions are recognized as competitive at the worldwide level for performance, energy efficiency, scalability and usability.
- The developed technologies are scalable and affordable not only for high-end centres, but for a broad set of industry segments.
- European HPC ecosystem differentiating through innovation leadership.
- European developed technologies represent a significant growing part of the added value of the world-wide HPC market and are relevant for other markets.

7. LINKS WITH OTHER INITIATIVES

The ETP action will align with the Horizon 2020 Program of the European Commission. Our intention is to be both a source of ideas for this program and a facilitator for its funding to achieve the ETP objectives. As an industrial forum, the ETP will present the point of view of European HPC technology providers and it will interact with the Commission taking part in HPC related actions in the Horizon 2020 development.

The ETP will take into consideration societal challenges defined as priorities by the Horizon 2020 program. Some of them could provide a good test-bed to show the impact of the research program that will be proposed by the ETP. Some synergies could be found in actions related to the societal challenges and projects related to HPC technology development.

The ETP acknowledges the importance of the implementation of the Key Enabling Technologies (KET) strategy and it should be noticed that the ETP's strategy for transforming scientific leadership into industrial successes is completely aligned with the High Level Group recommendations. Even if HPC is not mentioned in the most often listed KETs (nano-technology, nano-electronics, bio-technology, advanced materials, photonics), all these fields will heavily rely on HPC to progress. Thus, we consider HPC a 'Key Enabling Technology' for other KETs, or even KET itself.

PRACE's main objective is to develop a world-class HPC infrastructure for the European researcher. The ETP supports the strengthening of this initiative and we foresee a fruitful cooperation with PRACE. The expertise and skills developed by the ETP's program will help European researchers make the most effective use of current and future HPC systems. A strong and advanced user community will help the ETP to select relevant research directions that will impact the performance of real applications.

The ETP will also work closely with the existing HPC related projects such as EESI2 or PlanetHPC. The objective will be to have complementary actions and to avoid redundancy. The ETP will be the voice of the industry and will support the idea of European technology development. The ETP as a permanent structure can bring the continuity in the effort needed to give to HPC the place it deserves.

The ETP will take into account the existing works that are engaged within the Exascale initiative of the European Commission.

The ETP will develop its relationship with existing JTIs (Joint Technology Initiatives) such as Artemis (embedded systems) and ENIAC (nano-electronics). As mentioned previously, it is important both to see how HPC technologies could impact broader markets and to use as much as possible relevant technologies coming from other IT domains. The cooperation with these JTIs will be a way to find possible synergies. The ETP can also help these JTIs to understand the potential of HPC for their own objectives and so to increase the awareness and funding for HPC.

The same approach will be followed for the Eureka clusters such as ITEA2 and Catrene that deal with technologies connected to HPC.

The ETP will work at the level of the Member States when appropriate. The objective will be again to increase the awareness of the importance of HPC and to obtain support for the development of European HPC technology.

The ETP will establish a structured permanent link with end user industrial associations and SME industrial district in order to measure continuous improvement in the effectiveness of the technological development roadmap and to maintain a direct connection with the needs and requirements of the potential market segments represented by the industries which will benefit of HPC innovative solutions for a global European economical growth.

On a global level, ETP4HPC will foster cooperation with other regions by being active in initiatives such as the IESP (International Exascale Software Project) that gather the efforts of several regions to advance more rapidly toward a new generation of HPC systems. The ETP can represent the voice of the European players in this kind of forum and promote their expertise and capacities.

The ETP will be active in standardization initiatives and will support the emergence of standards that are important for European stakeholders.

8. CONCLUSIONS AND RECOMMENDATIONS

As shown in this document and in other reports, HPC is a key enabling technology impacting industrial competitiveness and research excellence. It is an important contributor to innovation and to societal challenge solutions. Being at the leading edge in HPC can mean having the ability to make a difference in an increasingly competitive world and also to design and propose worldwide solutions for sustainable development.

Getting a real leadership in HPC can be achieved by increasing the awareness about its importance, developing its use in existing and emerging domains, improving its impact and relevance in the education system, building a strong R&D capability and a strong HPC technology ecosystem. The latter is mandatory to develop excellence in HPC usage and to realize all the benefits of this domain for the European economy.

The organizations creating this ETP want to take a step in this direction. They believe that by leveraging the existing strengths of European actors, it is possible to improve the current situation of Europe in the HPC technology domain. They propose to establish a coordinating initiative led by the industry with this objective.

The first recommendation of the ETP4HPC is **to launch a research program with the aim to develop European technology in all the segments of the HPC solutions value chain.** The consequences of this action will be a strong HPC technology ecosystem meaning:

- Competitive HPC solutions designed in Europe.
- An economically sustainable HPC sector.
- Better ability to use HPC and make it a competitive advantage for European research and industry.
- Reduced reliance upon technology from other regions that are investing heavily in their HPC technology development.
- Competitiveness of existing industrial sector and development of new activities.
- Solutions for societal challenges

The second recommendation is **to focus** on the technical domains described in section 6.2 and **to prioritize** the following:

- Building on existing strengths.
- Analyzing disruptions that can change the current HPC landscape and facilitate introduction of new technologies developed in Europe.
- Selecting technologies with a market potential large enough for a sustainable development.
- Choosing technologies fitting the needs of important applications.
- Creating a favorable environment for SMEs (creation of start-ups and development of existing SMEs).
- Using synergies with other IT market technologies.

The ETP4HPC is dedicated to propose a Strategic Research Agenda (SRA) that will be designed along these lines and aims to influence the allocation of HPC research funds and investments in Europe. The SRA will define the research priorities for Europe's HPC industry and the ETP will facilitate the coordination of European, national and regional research programs in order to increase the return on investment of the efforts for the European stakeholders.