Investigation Report on Existing HPC Initiatives

CONTRACT NO EESI 261513
INSTRUMENT CSA (Support and Collaborative Action)
THEMATIC INFRASTRUCTURE

Due date of deliverable: 30.09.2010
Actual submission date: 29.09.2010
Publication date:

Start date of project: 1 JUNE 2010 Duration: 18 months

Name of lead contractor for this deliverable: EPSRC
Authors: Xu GUO (EPCC), Damien Lecarpentier (CSC), Per Oster (CSC), Mark Parsons (EPCC), Lorna Smith (EPCC)
Name of reviewers for this deliverable: Jean-Yves Berthou (EDF), Peter Michielse (NCF)

Abstract: This deliverable investigates large-scale HPC projects and initiatives around the globe with particular focus on Europe, Asia, and the US and trends concerning hardware, software, operation and support.

Revision JY Berthou, 28/09/2010

Project co-funded by the European Commission within the Seventh Framework Programme (FP7/2007-2013)		
Dissemination Level : PU		
PU	Public	X
PP	Restricted to other programme participants (including the Commission Services)	
RE	Restricted to a group specified by the consortium (including the Commission Services)	
CO	Confidential, only for members of the consortium (including the Commission Services)	

Table of Contents

Table of Figures

Glossary

Abbreviation / acronym	Description
ANR	Agence Nationale de la Recherche (France)
ASC	Advanced Simulation and Computing (US programme)
ASCI	Accelerated Strategic Computing Initiative (US)
ASCR	Advanced Scientific Computing Research (US programme)
ANL	Argonne National Laboratory (US)
ANURAG	Advanced Numerical Research and Analysis Group (India)
AWE	Atomic Weapons Establishment (UK)
BARC	Bhabha Atomic Research Center (India)
BSC	Barcelona Supercomputing Center (Spain)
CAE	Chinese Academy of Engineering (China)
CAS	Chinese Academy of Sciences (China)
CCRT	Centre de Calcul Recherche et Technologie (France)
C-DAC	Center for Development of Advance Computing (India)
CEA	Commissariat à l'Energie Atomique (France)
CINES	Centre Informatique National de l'Enseignement Supérieur (France)
CMCC	Centro Euro-Mediterraneo per i Cambiamenti Climatici (Italy)
CNRS	Centre National de la Recherche Scientifique (France)
CREST	Core Research of Evolutional Science & Technology (Japan)
CSCI	Comité Stratégique du Calcul Intensif (France)
CSIC	Consejo Superior de Investigaciones Cientificas (Spain)
CSTP	Council for Science and Technology Policy (Japan)
CSC	CSC – Finnish IT Center for Science (Finland)
DARPA	Defense Advanced Research Projects Agency (US)
DEISA	Distributed European Infrastructure for Supercomputing Applications (Europe)
DoD	Department of Defense (US)
DoE	Department of Energy (US)
EADI	Exascale Application and Data Initiative (Japan)
ECMWF	European Centre for Medium-Range Weather Forecasts (UK)
EDF	Electricité de France (France)
EESI	European Exascale Software Initiative (Europe)
EIC	Exascale Innovation Center (Germany)
EPCC	Edinburgh Parallel Computing Center (UK)
EPSRC	Engineering and Physical Sciences Research Council (UK)
ECL	Exacluster Laboratory (Germany)

FORTH	Foundation for research and Technology (Greece)
FZJ	Forschungszentrum Julich (Germany)
GCS	Gauss Centre for Supercomputing (Germany)
GENCI	Grand Equipment National de Calcul Intensif (France)
HEC-IWG	High End Computing Interagency Working Group (US)
HECRTF	High-End Computing Revitalization Task Force (US)
HLRS	High Performance Computing Center Stuttgart (Germany)
HPCI	High-Performance Computing Infrastructure (Japan)
HPCMP	High Performance Computing Modernization Program (US)
HSC	High Performance Computing Strategy Committee (UK)
IDC	International Data Corporation
IESP	International Exascale Software Project
IGST	Indian Grid for Science and Technology (India)
INCITE	Innovative and Novel Computational Impact on Theory and Experiment (US programme)
INRIA	Institut National de Recherche en Informatique et Automatique (France)
IRCSET	Irish Research Council for Science, Engineering and Technology (Ireland)
ISICLES	Ice Sheet Initiative for Climate Extremes (US)
JAEA	Japan Atomic Energy Agency (Japan)
JAMSTEC	Japan Agency for Marine-Earth Science and Technology (Japan)
JST	Japan Science and Technology Agency (Japan)
KAUST	King Abdullah University of Science and Technology
KISTI	Korean Institute of Science and technology Information (South Korea)
LANL	Los Alamos National Laboratory (US)
LBNL	Lawrence Berkeley National Laboratory (US)
LLNL	Lawrence Livermore National Laboratory (US)
LRZ	Leibniz Supercomputing Centre (Germany)
MAC	Munich Centre of Advanced Computing (Germany)
MSCS	Munich Computational Sciences Centre (Germany)
MEST	Ministry of Education, Science and Technology (South Korea)
METI	Ministry of Economy, Trade and Industry (Japan)
MEXT	Ministry of Education, Culture, Sports, Science and Technology (Japan)
MIC	Ministry for Internal Affairs and Communications (Japan)
MOST	Ministry of Science and Technology (China)
MSU	Moscow State University (Russia)
NAG	Numerical Algorithms Group (UK)
NAL	National Aerospace Laboratory (India)
NASA	National Aeronautics and Space Administration (US)

NCF	National Computing Facilities (The Netherlands)
NGS	Next Generation Supercomputer (Japan)
NICS	National Institute for Computational Sciences (US)
NITRD	Networking and Information Technology Research and Development (US programme)
NERSC	National Energy Research Scientific Computing Center (US)
NNSA	National Nuclear Security Administration (US)
NSCS	National Supercomputing Center in Shenzhen (China)
NSFC	National Natural Science Foundation of China (China)
NSF	National Science Foundation (US)
NSTC	National Science and Technology Council (South Korea)
NUDT	National University Defense Technology (China)
NUIG	National University of Ireland Galway (Ireland)
OHPC	Omnipresent High Performance Computing (US programme)
OPL	Open Petascale Libraries (Japan)
ORAP	Organisation Associative du Parallélisme (France)
ORNL	Oak Ridge National Laboratory (US)
OST	Office of Science and Technology (UK)
PERCS	Productive Easy-to-use Reliable Computing System (US)
PRACE	Partnership for Advanced Computing in Europe
PROSPECT	Promotion of Supercomputing Partnerships for European Competitiveness and Technology
RES	Red Espaňola de Supercomputación (Spain)
RUSNANO	Russian Corporation of Nanotechnologies (Russia)
RZG	Rechenzentrum Garching (Germany)
SciDAC	Scientific Discovery through Advanced Computing (US programme)
STFC	Science and Technology Facilities Council (UK)
STPI	Software Technology Parks of India (India)
TCD	Trinity College Dublin (Ireland)
TEXT	Towards Exascale Applications (Europe)
TGCC	Très Grand Centre de Calcul (France)
TUM	Technical University of Munich (Germany)
UCC	University College Cork (Ireland)
UHPC	Ubiquitous High Performance Computing (US program)
UVSQ	Université Versailles Saint Quentin-en-Yvelines (France)

1. Executive summary

This report provides details of large-scale HPC initiatives in the US, Asia and Europe and examines trends and future developments in this area. Current Petascale initiatives are described and future projects looking at Exascale development considered. The goal is to understand Europe's current position in the international HPC landscape, understanding our strengths and weaknesses.

The aim of the European Exascale Software Initiative (EESI) is to build a European vision and roadmap to address the challenges of developing and utilizing future large-scale HPC systems. This report contributes to this aim by providing a snapshot of the current state of the art in the area, providing a baseline and guide for this road mapping activity.

The main findings of the report follow.

- Following significant investment over the last few years, the US currently dominates the supercomputing sector, hosting two of the three Petascale systems[1] with significantly more expected by the end of 2010.
- To date, the US has led much of the discussion around Exascale computing, investing in the International Exascale Software Project (IESP) and funding programmes such as the Ubiquitous High Performance Computing program and the co-design centre call. However currently no multi-agency national Exascale initiative exists to drive the Exascale agenda forward.
- China has invested heavily in supercomputing technology and now holds second position in the TOP500 list. The Chinese government has actively promoted independent innovation to avoid reliance on foreign technology.
- Although China's recent top-end systems are impressive, it is notable that the peak performance / sustained performance ratio is high, leading to questions over how applicable these systems are.
- Japan and South Korea both expect to build Petaflop systems. Japan will have their first system by the end of 2010, and has a flagship project to develop a 10 Petaflop system by 2011-2012. South Korea has announced plans for a Petaflop system by 2013.
- Although India plans to have a Petaflop system by 2013, funding has recently been directed at small to medium range HPC centres.
- BULL is the only European HPC manufacturer. BULL will deliver two Petaflop systems in France respectively to CEA DAM, in 2010, and to GENCI in 2011. Up to now, Europe has not sought to develop its own "domestic technologies" such as in Asia. Europe has focussed on demonstrating expertise in software development and has recently seen a series of public-private partnerships between European research centres and international HPC producers to address the challenges associated with Exascale.
- Significant funding has been made available in Europe through the Framework 7 programme, both in Petascale (e.g. PRACE) and Exascale terms (e.g. the current Exascale computing, software and simulation call).
- The recent G8 Research Councils Initiative on Exascale applications for global issues is significant as it provides funding to enable collaboration between US and European partners and that it focuses on the applications of Exascale rather than hardware.
- Most Exascale initiatives across the continents identify a similar set of global issues as scientific drivers for Exascale computing. These include climate change, sustainable energy sources and nanotechnology.

[1] This figure is based on sustained, LINPACK, performance.

2. Introduction

In June 2010, the European Exascale Software Initiative (EESI) was launched with the aim of building a roadmap and set of recommendations to address the challenge of utilizing the current and next generation of supercomputers. Currently, top end systems have reached Petaflop performance and are anticipated to reach Exaflop performance by the end of the decade.

This report is a preliminary investigation of large-scale HPC projects and initiatives around the globe with particular focus on Europe, Asia, and the US. The target is to identify and map trends in hardware, software, operation and support. Information has been obtained mostly from public material, such as press releases and reports, but also from interviews with key persons in the various countries examined. The report has benefited from the knowledge and network of contacts available within the EESI consortium.

The majority of large-scale HPC initiatives and of HPC technology development has taken place in the US, Europe and Asia. Recently, significant progress has been made in the race to achieve a Petaflop, and today each of the three regions owns at least one system with Petaflop performance.

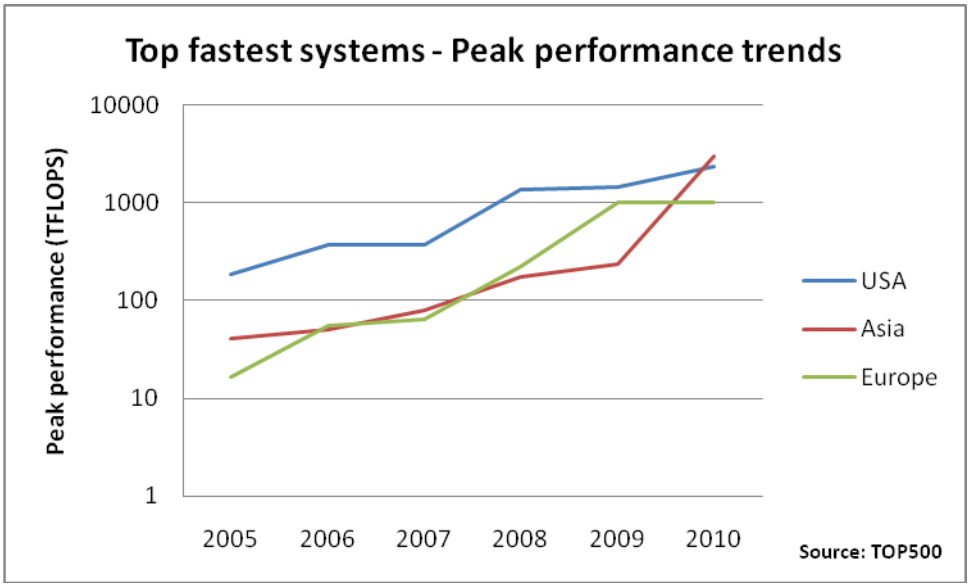

Figure 1: #1 systems (Peak performance) in the U.S., Asia, and Europe from 2005 to 2010 (Source TOP500 June lists 2005-2010)

China is the rising HPC country in Asia, with increasing number of systems listed in TOP500[2] and their fastest system now surpassing Japanese and European systems, and even the U.S. on peak performance. As of 2010, China is the only Asian country that has joined the "Petaflop club" (together with the U.S. and Europe (Germany)) although it should be followed soon by Japan, Russia and South Korea. In Europe, in 2009, Germany became the second country in the world to install a Petaflop

[2] http://www.top500.org/

supercomputer. Other European countries, boosted by the PRACE initiative, are likely to follow very soon, such as France in 2011.

In each of the three regions, large-scale initiatives have been indentified that are likely to support the migration from Petascale to Exascale systems, although through different paths. In addition to the national and regional initiatives covered by this report, it is important to note that major international initiatives targeting Exascale are also taking place. The growing hardware capabilities lead to outstanding technological breakthrough possibilities in computations and simulations which will be reached only through extensive international collaboration. This is the main rationale behind the creation of the International Exascale Software Project (IESP) and the European Exascale Software Initiative (EESI). We are also witnessing the launch of joint international calls targeting Exascale development, such as the recent G8 Research Councils Initiative on Multilateral Research Funding call. Announced in February 2010, this call proposals is for a new software programme called the "Interdisciplinary Programme on Application Software towards Exascale Computing for Global Scale Issues" with 10 M€ funding.[3] These initiatives will undoubtedly boost the migration towards Exascale systems.

Section 3.1 summarizes the HPC situation in the US, Section 3.2 the situation in Asia and finally the situation in Europe is described in Section 3.3.

[3] http://www.dfg.de/en/research_funding/international_cooperation/research_collaboration/g8-initiative/

3. HPC large scale initiatives in the U.S., Asia and Europe

3.1 HPC initiatives in the U.S.

Population: 307 M
GDP: 14,270,000 $M[4]
Yearly average supercomputer spending: 1276.1 $M[5]
Supercomputer spending / GDP: 0,0089%[6]

3.1.1 General overview

The U.S. has dominated the Supercomputing market for some considerable time. Since the TOP500 list began in 1993, the U.S. has had the dominate share of systems and peak performance. In 2002 however the Japanese Earth Simulator became the world's fastest supercomputer, leading to questions about the status of high-end computing in the United States. This, coupled with a realisation that the rapid growth in commodity based clusters threatened U.S. dominance in the HPC market and in their superiority in national security applications led to significant change and investment in the industry[7]. In 2002 the High-End Computing Revitalization Task Force (HECRTF) was formed to develop a five year plan for future investment in HPC[8]. In addition, the Advanced Scientific Computing Advisory Committee published a report detailing the "High-Performance Computational Needs and Capabilities in the Office of Science"[9]. This was re-enforced in the 2005 State of the Union address which indentified Supercomputing as important in maintaining competitiveness in the US. Since then, significant investment has been made in high-end systems. More recently, President Obama's administration listed exascale computing as one of the Grand Challenges of the 21st Century in their "Strategy for American Innovation"[10] and explicitly called for an "exascale supercomputer capable of a million trillion calculations per second".

3.1.2 HPC funding bodies in the U.S.

The main consumers of Supercomputing technology are the Department of Energy (DoE), the Department of Defense (DoD) including the Defense Advanced Research Projects Agency (DARPA), the National Science Foundation (NSF) and the National Aeronautics and Space Administration (NASA). The Networking and Information Technology Research and Development (NITRD) Program (HPC Act P.L. 102-194)[11] coordinates these organisations providing a framework and strategy for the support of US research and development in networking and information technology. High end

[4] Source: IDC report 2010 (The development of a Supercomputing Strategy in Europe" (SMART 2009/005, Contract No 2009/S99-142914)
[5] Over the last five years. Source: IDC report 2010
[6] Source: IDC report 2010
[7] DARPA's HPCS Program: History, Models, Tools, Languages", Dongarra et al, See: http://icl.cs.utk.edu/news_pub/submissions/final-1120.pdf (Accessed Aug 2010).
[8] The Roadmap for the Revitalisation of High-End Computing, ed D. Reid, 2003, See: http://www.cra.org/uploads/documents/resources/rissues/supercomputing.web_.pdf (Accessed Aug 2010).
[9] High-Performance Computational Needs and Capabilities in the Office of Science, 2004. See: http://www.er.doe.gov/ascr/ASCAC/Reports/FinalReport%20-HPC%20Needs.pdf (Accessed Aug 2010).
[10] A Strategy for American Innovation: Driving Towards Sustainable Growth and Quality Jobs, 2009 http://www.whitehouse.gov/administration/eop/nec/StrategyforAmericanInnovation/ (Accessed Aug 2010).
[11] See: http://www.nitrd.gov/

computing is coordinated by the Interagency Working Group on High End Computing (HEC IWG). NITRD's strategic plan for 2010 is currently available[12].

Within the DoE, the Advanced Simulation and Computing (ASC) Program, formally the Accelerated Strategic Computing Initiative (ASCI), provides simulation capability to "analyze and predict the performance, safety and reliability of nuclear weapons and to certify their functionality"[13] and supports the National Nuclear Security Administration (NNSA) Defense Program. This program supports Los Alamos National Laboratory (LANL), Lawrence Livermore National Laboratory (LLNL) and Sandia National Laboratories.

The DoE's Advanced Scientific Computing Research (ASCR) program aims to develop the algorithms, computer programs and hardware that advance scientific research in the DoE's remit. ASCR oversees facilities such as the National Energy Research Scientific Computing Center (NERSC) which is based at Lawrence Berkeley National Laboratory (LBNL) in California, the Leadership Computing Facility based at Oak Ridge National Laboratory's (ORNL) Center for Computational Sciences in Tennessee and Argonne National Laboratory (ANL) in Illinois. A draft of ASCR's strategic plan for high performance computing is now available, providing details of their plan for HPC resources over the next ten years[14].

The DoD also invests heavily in supercomputing technology, through DARPA (for example, through the DARAPA HPCS program, see below) and through the High Performance Computing Modernization Program (HPCMP). HPCMP was initiated in 1992 and aims to deliver world-class high-end facilities to the DoD's science and technology and test and evaluation communities. DARPA aims to "maintain the technological superiority of the U.S. military and prevent technological surprise from harming our national security by sponsoring revolutionary, high-payoff research bridging the gap between fundamental discoveries and their military use"[15]. Exascale computing falls within this remit.

The National Aeronautics and Space Administration (NASA) provides HPC resources for NASA scientists at it Advanced Supercomputing Division (ASD), providing supercomputing resources and simulation tools to carry out critical NASA missions. NASA's mission is to pioneer the future in space exploration, scientific discovery and aeronautics research and is known worldwide. Finally the National Science Foundation (NSF) funds supercomputers for open scientific research, with scientific areas such as climate change, clean energy, biology, and nanotechnology benefitting.

[12] See: http://www.nitrd.gov/DraftStrategicPlan/NITRDstratplan_Public_Comment.pdf and http://www.nitrd.gov/pubs/2011supplement/FY11NITRDSupp-FINAL-Web.pdf (accessed Sept 2010

[13] See: http://www.sandia.gov/NNSA/ASC/ (Accessed Aug 2010)

[14] Delivering Computing for the Frontiers of Science, Facilities Division Strategic Plan, 2010, See: http://www.er.doe.gov/ascr/ProgramDocuments/Docs/ASCRFacilitiesStrategicPlan.pdf (Accessed Aug 2010).

[15] http://www.darpa.mil/mission.html

Organisation	Office/Directorate/ Agency	Program / Project	Labs/ Centres Supported	Information
DoE	NNSA	ASC	LANL, LLNL, Sandia	Performance and reliability of nuclear weapons
DoE	Office of Science	ASCR	ORNL, ANL, NERSC (at LBNL)	Advance scientific research in DoE's remit
DoD	Office of the Security of Defense	HPCMP	AFRL, ARSC, ARL, ERDC, NAVY, MHPCC[16]	Accelerate development and transition of advanced defence technologies into superior war fighting capabilities
DoD	DARPA	HPCS, UHPC, OHPC[17]	Project funded	Maintain technological superiority of the US military
NASA	Exploration Technology Directorate, Ames Research Centre	High End Computing Capability project Project (HECC)	NASA Advanced Supercomputing Division (NAS)	Support critical NASA missions in e.g. space exploration
NSF	Office of Cyberinfrastructure		Project funded	open scientific research, including e.g. climate change, clean energy and biology

3.1.3 Current Supercomputing Systems

Seven of the systems in the top 10 from the TOP500 list are from the United States. At number one is the Jaguar system at Oak Ridge National Laboratory. Funded by the Department of Energy, by the ASCR program, this is a Cray XT system. Starting as a 26 TFLOPS machine in 2005 a $200 Million contract provided for a series of upgrades through to 2009. In November 2009 the system was upgraded to a 2.6 PFLOPS system (peak performance), through a $19.9 Million fund from the American Recovery and Reinvestment Act 2009 and is currently two systems: a Cray XT5 system with 37,376 AMD six-core Istanbul 2.6 GHz processors and a Cray XT4 system with 7,832 AMD four-core Budapest 2.1 GHz processors. The system is used for a wide variety of science, through the INCITE program (described below) and include fusion, climate change, astrophysics, new materials and protein engineering.

Project name	**Jaguar**
Short description	2.6 PFLOPS system, Number 1 on the TOP500 list
Url	http://www.nccs.gov/computing-resources/jaguar/
Organisations involved	Oak Ridge National Laboratory
Funding	Funded by the DoE ASCR program, for $200 Million in 2005 to upgrade the current XT3, further $19.9 Million in 2009.

[16] Air Force Research Laboratory (AFRL DSRC), Arctic Region Supercomputing Center (ARSC DSRC), US Army Research Laboratory (ARL DSRC), US Army Engineer Research and Development Center (ERDC DSRC), Navy DoD Supercomputing Resource Center (NAVY DSRC), Maui High Performance Computing Center (MHPCC DSRC)

[17] High Productivity Computing Systems Program (HPCS), Ubiquitous High Performance Computing program (UHPC), Omnipresent High Performance Computing (OHPC) program

Schedule	2005-2009
Application field	Various: fusion, protein engineering, climate change, new materials, astrophysics.

The third fastest machine in the world is Roadrunner, an IBM hybrid system with 12,960 IBM PowerXCell 8i and 6,480 AMD Opteron dual-core processors and a peak performance of 1.38 PFLOPS. It is funded by the DoE's NNSA for $125M and is being used to simulate how nuclear material ages to predict whether aging nuclear weapons are safe and reliable. Coming into operation in 2008 it has been in development since 2002.

Project name	**Roadrunner**
Short description	1.38 PFLOPS system, Number 3 on the TOP500 list
Url	http://www.lanl.gov/roadrunner/
Organisations involved	Los Alamos National Laboratory
Funding	$900M, $125M
Schedule	2002-2008
Application field	Nuclear material aging

The fourth fastest machine in the world is Kraken, an NSF funded system, with $65M awarded to the National Institute for Computational Sciences (NICS) at the University of Tennessee. The system is a Cray XT5 with AMD Istanbul Opteron six core 2.6 GHz processors and has a peak performance of 1.02 PFLOPS. Installed in 2009 it is used by various scientific fields including climate, fusion energy, biology, lattice QCD, and astrophysics.

Project name	**Kraken**
Short description	1.02 PFLOPS system, Number 4 on the TOP500 list
Url	http://nics.tennessee.edu/
Organisations involved	National Institute for Computational Sciences (NICS) at the University of Tennessee
Funding	$65M
Schedule	2009
Application field	Climate, fusion energy, biology, lattice QCD, and astrophysics

Other significant U.S. systems in the top 10 machines in the TOP500 list include: Pleiades, an SGI Altix ICE system based at NASA Ames Research Centre; the BlueGene/L system at LLNL; Intrepid, the BlueGene/P system at Argonne National Labs and Red Sky, the Sun Blade system at Sandia National Laboratories.

3.1.4 Supercomputing Applications Areas and Funding

The Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program allocates 60% of the run-time on the DoE's Leadership class facilities (Jaguar at ORNL and the BlueGene/P system at ANL). The program is open to all researchers from both academia and industry and awards sizeable chunks of CPU time for a variety of grand challenge problems. Started in 2003 it has been running for 7 years. Projects have ranged from astrophysics to nanotechnology and focus on new forms of energy and climate change. In 2010, 1.6 billion supercomputing processor hours were allocated to 69 cutting-edge research projects and in 2011 roughly the same will be awarded again (expected awards will be announced in November 2010).

The ASCR leadership challenge allocated over 610M processor hours in 2010 on the resources at NERSC and the Leadership Computing Facilities at Argonne and Oak Ridge (this equates to up to 30% of these resources). The program funds researchers for projects of special interest to the DoE's

energy mission "*in areas such as advancing the clean energy agenda and understanding the Earth's climate, for national emergencies, or for broadening the community of researchers capable of using leadership computing resources*".[18]

The Scientific Discovery through Advanced Computing (SciDAC) program started in 2001 as a five year project from the DoE to develop software and hardware infrastructure for terascale systems. This has since become a highly successful program in the U.S. and aims to fund collaborative research projects between interdisciplinary teams working on major software and algorithm development for applications to problems in the SC core programs (namely Basic Energy Sciences, High Energy Physics, Nuclear Physics, Advanced Scientific Computing Research, Fusion Energy Sciences, and Biological and Environmental Research). For example the ASCR SciDAC Initiative recently funded a series of projects, termed ISICLES (Ice Sheet Initiative for CLimate ExtremeS) to accelerate research for algorithms and codes for ice sheet models that are scalable at extreme scales of computing[19].

3.1.5 Current funding and near Future Petascale Systems in the U.S.

The Defense Advanced Research Projects Agency (DARPA) approved The High Productivity Computing Systems Program (HPCS) in 2001, a program aimed at providing economically viable high productivity systems, which are substantially easier to program and use than existing systems, by the end of 2010. Carried out over three phases, this program still funds two major petascale initiatives, one with Cray and one with IBM. During Phase 3, IBM was awarded $244 Million to develop a prototype PERCS (Productive Easy-to-use, Reliable Computing System) by 2010. Since then NSF have invested $208 Million over four years (from 2007) to produce the first production deployment of PERCS technologies at University of Illinois at Urbana-Champaign, called Blue Waters (note this is a different machine from the DARPA system). Also in Phase 3, Cray was initially awarded $250 Million[20] to develop a prototype based on its Cascade hybrid architecture. This builds on a Phase 2 investment of $43.1M for Cray and $53.3M for IBM. In addition to hardware, the HPCS program funded the development of the HPC Challenge benchmark suite, to help to define the performance boundaries of Petascale systems[21].

Blue Waters is expected to be one of the most powerful supercomputers in the world when it comes online in 2011. The system will be based on the IBM POWER7 hardware. It will have a peak performance of 10 PFLOPS (10 quadrillion calculations every second) and will achieve sustained performance of 1 PFLOPS running a range of science and engineering codes. The system is being developed by the University of Illinois at Urbana-Champaign, its National Center for Supercomputing Applications, IBM, and the Great Lakes Consortium for Petascale Computation. Announced in August 2008 and supported by a $208 million grant from the National Science Foundation the system is estimated to come online in 2011.

Cray have recently (April 2010) been awarded a $45M contract[22] from the NNSA to build a next generation Cray Supercomputer named Ceilo. Ceilo is expected to be installed in the second half of 2010 at LANL with upgrades expected in 2011 (a 1-2 Petaflop system), and will be used, for example, to develop new nuclear weapons and simulate their explosions. The system is expected to be a Petaflop computer and be a Cray XT/XE system.

IBM are currently developing Sequoia, a Petascale system to be housed at LLNL for the NNSA as part of the ASC program, scheduled to go on-line by 2011 with a target peak performance of 20 PFLOPS and based on the BlueGene/Q design. The system will be mainly used for nuclear testing but will also be available for scientific research such as climate change, astronomy and energy.

[18] See: http://www.sc.doe.gov/ascr/Facilities/ALCC.html Accessed Aug 2010.
[19] See: http://www.scidacreview.org/1002/html/news6.html (Accessed Aug 2010)
[20] This has since been scaled back by 60 M$.
[21] See: http://icl.cs.utk.edu/hpcc/ (Accessed Aug 2010)
[22] See: http://investors.cray.com/phoenix.zhtml?c=98390&p=irol-newsArticle&ID=1409130&highlight= (Accessed Aug 2010)

In May 2010 Cray announced that it had been awarded a $47M contract to build a Petaflop system for advanced climate modelling. The contract from the DoE will provide the National Oceanic and Atmospheric Administration (NOAA), which is overseen by the Department of Commerce (DoC) and the Oak Ridge National Laboratory (ORNL) with a Cray XT6 system in late 2010 (to be upgraded to an XE6 in 2011). The system is funded under the American Recovery and Reinvestment Act of 2009.

Finally Pleiades, the SGI system at NASA, Ames Research Centre is due to be upgraded to a Petascale system. Currently it has a performance of 970 TFLOPS (number 6 in the TOP500) but is due to be upgraded to 10 PFLOPS by 2012.

Program	Funder	Dates	Organisation	Award
HPCS	DARPA	2001-2010	IBM	$53.3M (Ph2)+$244M (Ph3)
HPCS	DARPA	2001-2010	Cray	$43.1M (Ph2) + $250M (Ph3)
BlueWaters	NSF	2008-2011	University of Illinois	$208M
Ceilo	NNSA	End 2010, upgrade 2011	LANL	$45M
Sequoia	NNSA	Scheduled 2011	LLNL	
Climate	NOAA	End 2010, upgrade 2011	ORNL	$47M
Pleiades	NASA	Upgrade 2012	Ames Research Centre	

3.1.6 Funded Exascale Projects in the U.S.

The International Exascale Software Project (IESP) is an international initiative to define a comprehensive roadmap for Exascale software. This is funded by the DoE and the NSF[23] and has recently published Version 1.0 of the roadmap[24].

DARPA have recently announced the four groups selected for the Ubiquitous High Performance Computing program (UHPC), a program aimed at developing " an innovative, revolutionary new generation of computing systems that overcomes the limitations of current evolutionary approach"[25] Prototype systems are expected by 2018, the initial awards are for four years with two follow on phases expected. "The resulting UHPC capabilities will provide at least 50-times greater energy, computing and productivity efficiency". The total award for the four groups chosen was $76.6M[26]. NVIDIA, Cray Inc and six University partners have been awarded a four year contract of $25 Million. Intel's project is worth $49 Million with equal investment from DARPA and Intel and involving SGI. Sandia National Labs together with industry partners Micron Technology, Inc. and LexisNexis Special Services[27], Georgia Institute of Technology, Atlanta, was selected to lead an Applications, Benchmarks and Metrics team for evaluating the UHPC systems under development[28].

DoE's ASCR program has a recent call for Exascale Co-design centres, "a computer system design process where scientific problem requirements influence architecture design and technology and constraints inform formulation and design of algorithms and software". Proposals were required by June 2010 and a decision is still awaited. Currently $8M is available for fiscal year 2010, with further years support dependant on fund availability. However they have requested projects of between $5M and $10M per year, for five years, with the possibility of renewing for a further 5 years.

[23] The value of this funding is unknown.
[24] International Exascale Software Project Roadmap; Dongarra, Beckman, et al., May 30, 2010, CS Technical Report: ut-cs-10-654, See: http://www.exascale.org/mediawiki/images/4/42/IESP-roadmap-1.0.pdf. Accessed Aug 2010.
[25] See: http://www.darpa.mil/news/2010/UHPCNewsRelease.pdf. (Accessed Aug 2010).
[26] See: http://www.theregister.co.uk/2010/08/09/darpa_extremescale_awards/ (Accessed Aug 2010)
[27] Funding value not known.
[28] Funding value not known.

In January 2010, the DoE's ASCR program also announced a call for proposals for "Advanced Architectures and Critical Technologies for Exascale Computing". This called for interdisciplinary teams to focus on addressing key barriers on the path to Exascale. With $5M available the program expected to fund around 5 projects. One of the successful projects was from Penn State, along with collaborators from Oak Ridge National Lab, Hewlett-Packard (HP), and the University of Michigan, for a $3M, 3-year project.

DARPA recently released a Broad Agency Announcement for the Omnipresent High Performance Computing (OHPC) program (initial closing date is in August 2010). The aim is to fund research that will dramatically advance the performance and capabilities of future systems and enable extreme scale computing. Proposals should provide significant advantages for the activities being performed under the UHPC program. DARPA expect to fund multiple projects of up to $1M a year, with the projects running for three years.[29]

Program	Funder	Dates	Organisation	Award
UHPC	DARPA	Prototype 2018	NVIDIA, Intel, Sandia, Georgia institute of Technology	$76.6M
Exascale Co-design	DoE (ASCR)	Call out 2010, 5 year projects, possible renewal for 5 years	To be decided	$8M in 2010, further expected
Advanced Architectures	DoE (ASCR)	2010-2013	Penn State, others?	$5M
OHPC	DARPA	Closes Aug 2010, 3 year project	To be decided	Multiple projects, $3M per project

3.1.7 Laboratory and Research Institute Initiatives

Numerous studies and reports have been written and initiative published in the U.S. on Exascale Computing and are too numerous to list here. The following summary aims to provide a flavour of the activities underway in the U.S. and is not intended to be comprehensive.

ASCR ran a series of town hall meetings in 2007 to discuss the challenges and problems in energy, the environment and basic science in achieving Exascale. A series of workshops then followed and identified Exascale requirements for energy research, environmental sustainability, national security, biology, astrophysics, climate science and nuclear physics. Reports from the various workshops are available on-line and provide details of grand challenge problems for Exascale.[30] In addition a set of DARPA reports summaries these results.[3132] The SciDAC review (issue sixteen) also provides a good description of grand challenge problems for Exascale[33] and provides a summary of goals of the DoE's Extreme Scale Initiative over the next 10 years. Finally Oak Ridge National Laboratory have published a report on "Science Prospects and Benefits with Exascale Computing"[34], which provides a good overview of scientific drivers for Exascale.

[29] See: http://insidehpc.com/2010/06/23/darpa-announces-omnipresent-hpc-program/ (Accessed Aug 2010).
[30] See: http://www.er.doe.gov/ascr/ProgramDocuments/ProgDocs.html (Accessed Aug 2010)
[31] See: http://www.er.doe.gov/ascr/Research/CS/DARPA%20exascale%20-%20hardware%20%282008%29.pdf (Accessed Aug 2010)
[32] See: http://users.ece.gatech.edu/~mrichard/ExascaleComputingStudyReports/ECSS%20report%20101909.pdf (Accessed Aug 2010)
[33] See: http://www.scidacreview.org/1001/index.html (Accessed Aug 2010)
[34] See: http://www.nccs.gov/wp-content/media/nccs_reports/Science%20Case%20_012808%20v3__final.pdf (Accessed Aug 2010)

Oak Ridge National Laboratory published a report on “Preparing for Exascale: ORNL Leadership Computing Facility Application Requirements and Strategy”, describing the fundamental science case driving the need for the next generation of computer hardware[35]. Argonne National Laboratory has a major initiative on ”Leadership Computing -- Exascale Technology for Scientific Computing” which aims to ” To develop scientific computing at the exascale in support of basic and applied science through innovations and breakthroughs in hardware design, software tools, and infrastructure.”[36] Sandia National Lab published a paper on “Software Challenges for Extreme Scale Computing: Going from Petascale to Exascale Systems ”discussing the issues of preparing applications for Exascale.[37] Finally a recent SOS (Sandia/Oak Ridge/Swiss Commodity-Based Computing Collaboration) workshop (SOS14, 8-11 March, 2010) focused on the challenges of Exascale systems.[38]

[35] See: http://www.nccs.gov/wp-content/media/nccs_reports/olcf-requirements.pdf (Accessed Aug 2010)
[36] See: http://www.anl.gov/Media_Center/strategicplan/strategic_plan.pdf (Accessed Aug 2010)
[37] See: http://www.sandia.gov/~maherou/docs/SoftwareChallengesForExtremeScaleComputing.pdf (Accessed Aug 2010)
[38] See: http://www.csm.ornl.gov/workshops/SOS14/ (Accessed Aug 2010)

3.2 HPC initiatives in Asia

3.2.1 General overview

Asia has been a leading region in the HPC field with countries such as Japan being at the forefront of HPC technology for many years. While Japan's commitment to pursue HPC leadership may have declined in recent years, the country continues to invest significant resources in this sector and has new upgrade plans that are likely to put it back into the HPC leadership race. However, China is the rising HPC country in the region, with increasing number of systems listed in TOP500 and China's fastest system now surpassing Japanese and European systems, and even the U.S. on peak performance. As of 2010, China is the only Asian country that has joined the "Petaflop club" (together with the U.S. and Germany) but it should be followed very soon by Japan. Other important HPC countries in the region to watch include South Korea and India[39].

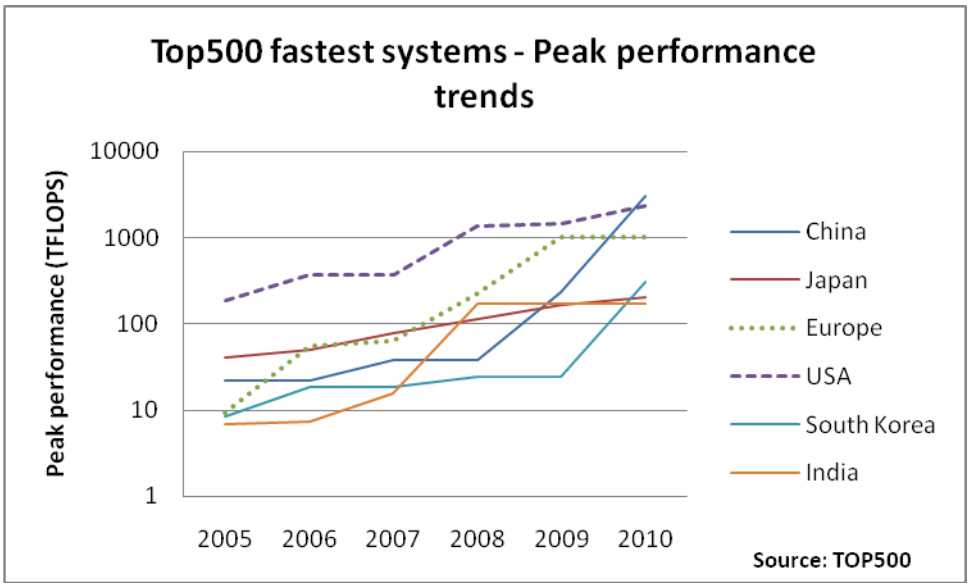

***Figure 2*: #1 systems (peak performance) in China, Japan, South Korea, India, compared with Europe and the U.S. from 2005 to 2010 (Source TOP500 June lists 2005-2010)**

The survey has identified a number of key HPC initiatives in these four countries that could potentially prepare the way for the exascale era. One general trend is that the countries in this region are also engaged in the development of "domestic technology", not only seeking to achieve the highest system performances but also to build their own systems and reduce their reliance on foreign technology.

It is also important to note that new plans related to HPC activities are likely to be unveiled in 2011, especially in Japan and China which are both preparing their next major policy framework regarding science and technology. The information we gathered should therefore be reviewed when these plans have been officially unveiled.

[39] Other interesting HPC countries include Hong-Kong, Singapore, and Taiwan. These are not examined in this preliminary survey. It can however be noted that Singapore and Hong-Kong, together with South Korea, are among the top four HPC-investing countries (ration of their supercomputing spending as compared to yearly GDP) in the world – the other top four nation is the US. Source: IDC Report 2010.

3.2.2 Japan

Population: 127 M
GDP: 5.049 $M[40]
Yearly average supercomputer spending: 278 $M[41]
Supercomputer spending / GDP: 0,0055%[42]

HPC policy and funding programs in Japan

Decision-making for science and technology occurs within the Council for Science and Technology Policy (CSTP), located within the Cabinet Office, and includes the direct involvement of the Prime Minister as well as Ministers from key offices of state such as the Ministry for Internal Affairs and Communications (MIC); the Ministry of Finance; the Ministry of Education, Culture, Sports, Science and Technology (MEXT); and the Ministry of Economy, Trade and Industry (METI)[43].

The Science and Technology Basic plans which are drawn up every five years in Japan drive most of Japan's HPC development activity. The current plan is the 3rd Science and Technology Basic Plan and covers the period 2006-2010.

The Japan Science and Technology Agency (JST) is responsible for implementing science and technology policies defined by the Science and Technology Basic Plan. The aims are to encourage the development of breakthrough technologies and to provide funding and support for researchers. CREST (Core Research of Evolutional Science & Technology) is one of JST's research promotion programs with an aim of stimulating achievement in fundamental science fields. JST distributes public funds from the MEXT to support a broad range of scientific research relevant to the nation's Strategic Sectors.

Currently, there is some political support in Japan for developing the infrastructure and significant effort is being made to organize the national HPC communities and institutions to ensure efficient use of computing resources such as its current flagship project developed by RIKEN, and to encourage the development of the HPC sector.

Current Supercomputing Systems in Japan

Japan has been for many years at the forefront of HPC technology with a large variety of architectures ranging from vector supercomputers to the specific gravity engine. While Japan's commitment to pursue HPC leadership may have declined in recent years, Japan's HPC infrastructure and ecosystem have remained very efficient and innovative.

In June 2010, 18 systems appeared in the TOP500 list. The fastest ranked #22 and was a Fujitsu cluster installed in 2009 at the Japan Atomic Energy Agency (JAEA) and featuring a theoretical peak performance of 200.08 TFLOPS (191.40 TFLOPS Linpack performance). Throughout the last decade Japan has regularly upgraded its systems and now hosts seven supercomputing centres with an average computing environment of 100 TFLOPS.

[40] Source: IDC report 2010
[41] Over the last five years. Source: IDC report 2010
[42] Source: IDC report 2010
[43] http://cordis.europa.eu/erawatch/index.cfm?fuseaction=ri.content&topicID=4&countryCode=JP

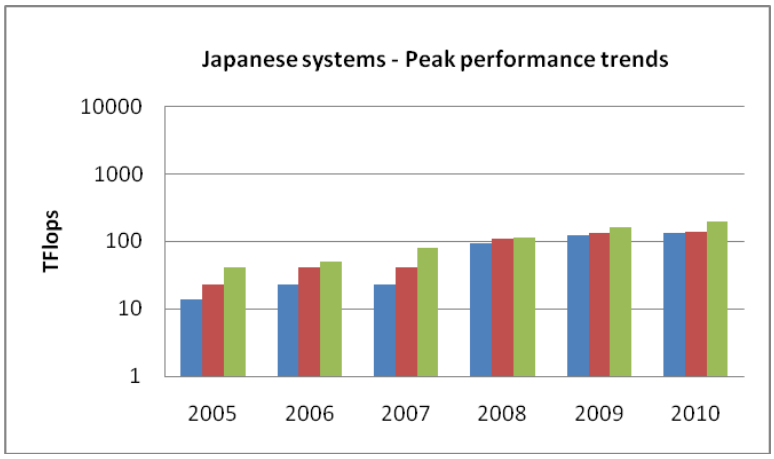

Figure 3*: *Peak performance of the three fastest Japanese supercomputers from 2005 to 2010 (source: TOP500 June lists)

Japanese universities and institutes are strong actors in the computing field, several of them hosting powerful systems in the 100 TFLOPS range, such as the University of Tokyo (Information Technology Center), the Tokyo Institute of Technology (GSCIC Center) and the University of Tsukuba (Center for Computational Sciences).

In addition, Japan has a number of manufacturers selling supercomputers such as Fujistu, NEC and Hitachi, these companies are leading providers of the national and international market, accounting for 11 out of 18 systems listed in TOP500.

Current Funding and near Petascale Systems in Japan

Although Japan has slowly slipped down in the TOP500 list – since June 2007, it has had no systems listed in the top 10 entries – new upgrade plans are likely to return Japan to a leadership role.

In June 2010, the Tokyo Institute of Technology announced the details of the "Tsubame 2.0," the next-generation supercomputer system for the Institute of Technology that is expected to start operation in November 2010[44]. Tsubame 2.0 will be Japan's first Petaflop supercomputer with computing capacity of 2.39 PFLOPS (30 times higher than its predecessor Tsubame 1.0). The system will be constructed by NEC Corp and Hewlett-Packard Co.[45]. The University also announced that it plans to introduce the Tsubame 3.0, a successor to the Tsubame 2.0, in 2014 or 2015 and development is already underway. The Tsubame 3.0 will aim for 30 PFLOPS with a power consumption equivalent or even less than of the Tsubame 2.0.[46]

However, the current flagship project in Japan is the RIKEN[47] Next Generation Supercomputer (NGS) Project that seeks to develop a 10 PFLOPS-class supercomputer in KOBE, about 450 Km of Tokyo. Around 1 billion Euros has been allocated to this project (covering the system, the facility, the applications, etc.) under the 3rd and Technology Basic Plan (2006-2010). The computer and research building was completed in May 2010, and the supercomputer is being installed and expected to run in 2011-2012.

[44] http://techon.nikkeibp.co.jp/english/NEWS_EN/20100618/183587/
[45] http://en.akihabaranews.com/51762/desktop/japans-first-Petaflop-supercomputer-tsubame-2-0-to-be-launched-this-autumn
[46] http://techon.nikkeibp.co.jp/english/NEWS_EN/20100618/183587/?P=3
[47] The Institute of Physical and Chemical Research (RIKEN) is a natural sciences research institute in Japan conducting research in many areas of science, including physics, engineering, chemistry, nanotechnology, biology, medical science, and computational science. It is funded by the Japanese government, and its annual budget in 2009 was approximately 1 billion US$ (around 800 M€). Source: http://www.riken.jp/engn/index.html

In terms of system configuration, RIKEN is a massive parallel system developed by Fujitsu and designed to meet diverse application environments. The system will have a sustained Petaflop performance on real applications.

A number of major applications have been selected for the Next-generation supercomputer in eight areas, such as manufacturing, disaster prevention, aerospace and the environment. Nanotechnology and Life sciences are targeted as “grand challenges”:

Project name	**RIKEN Next Generation Supercomputer**
Short description	Development of a 10 PFLOPS-class supercomputer
Url	http://www.nsc.riken.jp/index-eng.html
Organizations involved	RIKEN Next-Generation Supercomputer R&D Center + several partners industries, universities and government
Main contact person	Mr. Tadashi Watanabe, Project Leader, RIKEN Next-Generation Supercomputer R&D Center
Funding	Around 1 billion Euros has been allocated to this project (covering the system, the facility, the applications, etc.)
Schedule	2006-2012
Application field	Natural sciences, nanoscience, life sciences, engineering, disaster prevention, etc.

According to the current plan, the Advanced Institute for Computational Science (AICS) – which is being established in KOBE – will have a role in running the Next Generation Supercomputer as the centre for computational science in Japan. This new organization will run the Next Generation Supercomputer efficiently for users from a wide range of research areas, lead collaborations between computational and computer scientists (which is seen as key for building future exascale systems), and promote industrial use of HPC. The centre will also propose the future direction of HPC in Japan and create future hardware concepts for HPC after the Next Generation Supercomputer.

In order to develop the strategic use of RIKEN, the Japanese government has selected five strategic fields in science and technology to exploit the system. The program for each field is managed by a core institute responsible for organizing research and supercomputer resources for the community. Each core institute will receive priority allocation of NGS resources and funds to achieve the research goals:

- Life science/Drug design (Toshio YANAGIDA - RIKEN)
- New materials/energy creation (Shinji TUNEYUKI - University of Tokyo)
- Global climate change prediction for disaster prevention / mitigation (Shiro IMAWAKI - JAMSTEC)
- Manufacturing technologies (Chisachi KATO - University of Tokyo)
- The origin of matters and the universe (Shinya AOKI - University of Tsukuba)

In 2009, the goal of the NGS was reconsidered by the Japanese government who decided to put more emphasis on the user point of view. This led to a new plan for the creation of an Innovative High-Performance Computing Infrastructure (HPCI) consisting of various HPC sites including NGS and other universities and institutes supercomputing sites. The role of HPCI will be to establish the hierarchical organization of the Next-Generation Supercomputer with other supercomputers, to set up a large-scale storage system for the NGS and other supercomputers, and to establish an HPC consortium which will run the HPCI.

Towards Exascale computing

Although there is no clear published roadmap towards Exascale in Japan at the moment, the issue is being discussed and funding programmes and projects may soon take shape. The 4th Science and Technology Basic Plan, running for the period 2011-2015, is still under discussion and is likely to include plans to develop Exaflop class HPC technology[48].

The Japan Science and Technology Agency, which implements science and technology policies formulated by the Science and Technology Basic Plan, also has a number of funding programmes that could prepare the way towards Exascale, such as the JST CREST (Core Research for Evolutional Science and Technology) for "new HPC technologies" program which is to be software-oriented and have an overall budget of around 5M€ for 5-7 years[49]. In terms of software, research in Japan mostly focuses on programming and languages environments, numerical libraries, communication libraries and applications.

There are also a number of public-private initiatives taking shape that focus on software development towards Exascale. One of them is the Exascale Application and Data Initiative (EADI) led by Fujitsu which is in the final stages of developing a 10 PFLOPS system (the RIKEN NGS) for the Japanese Government. EADI is an umbrella project that collects together all of Fujitsu's activities towards Exascale computing.

This includes the Open Petascale Libraries project (led by Fujitsu Laboratories of Europe) which seeks to develop suitable thread-parallel and hybrid numerical libraries. The aim of the Open Petascale Libraries project is to produce numerical software suitable for the Fujitsu platform and similar massively parallel multi-core systems, and to make the libraries freely available to the computational science community as open-source software. The OPL Advisory Panel will consist of Professors Jack Dongarra, Bill Gropp, Kimihiko Hirao and Anne Trefethen, and the initial members are likely to be Fujitsu, the Innovative Computing Laboratory at the University of Tennessee, the Oxford e-Research Centre, the UK Science and Technology Facilities Council, Imperial College London, and the Numerical Algorithms Group Ltd. Projects are expected to range from those targeted at the immediate need for efficient thread-parallel and hybrid MPI/thread numerical libraries to projects that are more future-looking and that tackle some of the issues that have been identified in the IESP Roadmap. The consortium agreement between the different parties is under preparation and the initial membership and legal framework were expected to be in place by the end of August 2010, with a first Open Petascale Libraries Workshop to be held during SC10 in New Orleans.

Project	**EADI/Open Petascale Libraries (OPL)**
Short description	To produce numerical software suitable for the Fujitsu platform and similar massively parallel multi-core systems, and to make the libraries freely available to the computational science community as open-source software
Organizations involved	Fujitsu, the Innovative Computing Laboratory at the University of Tennessee, the Oxford e-Research Centre, the UK Science and Technology Facilities Council, Imperial College London, and the Numerical Algorithms Group Ltd
Main contact (for Europe)	Ross Nobes, Environment & Health Research Division Fujitsu Laboratories of Europe
Schedule	2010-

[48] http://www.prace-project.eu/documents/deisa-prace-symposium-2010-presentations/4_prace-deisa2010-ukawa-final.pdf
[49] idem

3.2.3 China

Population: 1.300 M
GDP: 4.758 $M[50]
Yearly average supercomputer spending: 67 $M[51]
Supercomputer spending / GDP: 0,0014%[52]

HPC Policy and funding programs in China

China has a highly centralized research system organised and controlled by the central government. The National Steering Group for Science and Technology and Education in the State Council coordinates all education, research and innovation related activities. The Ministry of Science and Technology (MOST), formerly known as the State Science and Technology Commission, is the leading ministry and works with other ministries or agencies, including the National Natural Science Foundation of China (NSFC), the Chinese Academy of Sciences (CAS), and the Chinese Academy of Engineering (CAE), to coordinate science and technology activities. NSFC develops science and technology programs and provides funding for basic and some applied research; CAS comprises high-level research institutes and together with CAE has a number of academic divisions of science and engineering.

Science and technology in China has been growing rapidly. As the country develops, the government has placed an increasing emphasis on science and technology, and strives to promote independent innovation in order to reduce the level of reliance on foreign technology. Every five years, the Chinese central government prepares a major plan setting direction in all key areas, including science and technology. As of 1990 this has included HPC research activities. The 11th five-year plan, which began in 2006 and will run until the end of 2010 has focused on the development of Petaflop supercomputers and a grid computing environment[53].

Since the 1980s, China has also formulated a series of national programs for science and technology research and development – often integrated to the five-year plans – with the strategic aim of improving China's competitiveness in science and technology in the 21st century. The National High Technology R&D Programme[54] (also called “863 Programme”) is often cited as an efficient tool for the implementation of China’s HPC policy, allowing sustained national investments in innovation and manufacturing. This sponsors research in key high technology fields, including Information Technology which was identified in 2002 as a priority research field[55].

Supercomputing environment

In just a few years, China has grown from a mid-range HPC country to an international leading one. China has kept increasing its number of systems and is now tied with Germany (steadily declining) in the top500, after the US, UK and France. In 2009, China became the third country to build a Petaflop supercomputer (Tianhe-I, a 1.206 peak PFLOPS machine installed at the National University Defense Technology (NUDT) in Changsha) and now has three systems with a peak performance above 1 Petaflop, two of which are ranked among the ten most powerful supercomputers in the world. In June 2010, a Chinese system called Nebulae, developed by Dawning and based in the National Supercomputing Center in Shenzhen (NSCS) became the world’s fastest computer with a theoretical

[50] Source: IDC report 2010
[51] Over the last five years. Source: IDC report 2010
[52] Source: IDC report 2010
[53] Dr. Kai Lu, The *Developments of HPC Market & Research in China*, presentation at the International Supercomputing Conference (ISC’10), Hamburg, 1st June 2010.
[54] http://www.most.gov.cn/eng/programmes1/200610/t20061009_36225.htm
[55] More information on the structure of the Chinese research system as well as on the different science & technology programmes can be found on the European portal ERAWATCH: http://cordis.europa.eu/erawatch/index.cfm

peak performance at 2.98 Petaflops (# 2 with a Linpack performance of 1.271 Petaflops). This is the highest ranking a Chinese system has ever achieved.

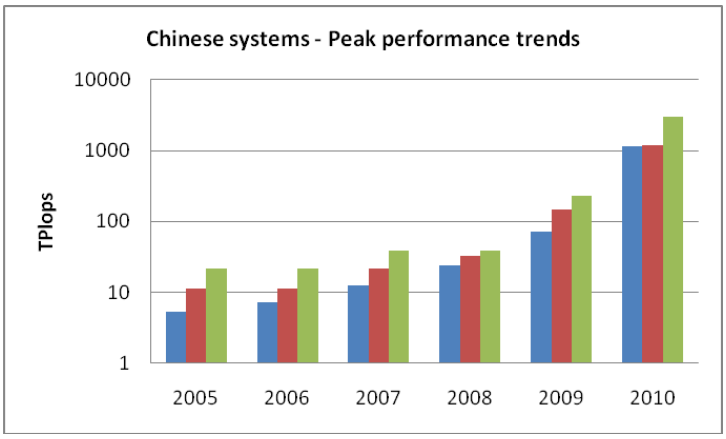

Figure 4*: *Peak performance of the three fastest Chinese systems from 2005 to 2010 (source: TOP500 June lists)

Although the peak performance of Chinese systems is very impressive, their Linpack performance is relatively low and does not achieve 50% efficiency (while most supercomputers have a Linpack efficiency of 75% or above). This may therefore limit their applicability. However even with this relatively low sustained performance, the three fastest Chinese supercomputers still rank highly in the top500 list (respectively #2, #7, and #19 in the TOP500 list of June 2010).

Another example of the commitment of the Chinese government to developing petascale technology is the creation in 2009 of China's first Petaflop supercomputing centre, the National Supercomputing Center in Shenzhen (NSCS), where the Dawning Nebulae machine is installed. Sponsored by the Ministry of Science and Technology (MOST), Chinese Academy of Science (CAS) and Shenzhen government this was created with an initial 90 M€ investment[56]. NSCS is expected to act as a hub for research and innovation in China and to serve as a platform for science and technology, industrial innovation, research and development and training. The centre will facilitate the implementation of "national S&T projects, industrial innovation, digitized urban management, and basic research".[57]

China has a number of growing HPC manufacturers such as Dawning, Lenovo, the Institute of Computer in the National University of Defense Technology (NUDT), Inspur, and Sunwei. These manufacturers are steadily increasing their share in the national market, providing 45% of China's top100 supercomputers in 2009, with Dawning taking the biggest share (27%), ahead of IBM (26%) and HP (23%)[58]. This is often viewed as a strategic issue for China's independence. These companies are also eager to compete further on the international HPC market.

HPC in China has developed rapidly in recent years, especially in HPC hardware which now competes on the world stage. There is however still a gap in terms of usage and applications. Prof. Wenguang Chen of Tsinghua University recently stated in an interview with CNETNews that "There are only a few of large-scale parallel applications at this stage and further developments are required"[59].

[56] http://english.peopledaily.com.cn/90001/90776/90881/6815913.html
[57] http://www.most.gov.cn/eng/newsletters/2009/200906/t20090609_71192.htm
[58] China Top100
[59] http://www.cnetnews.com.cn/2010/0716/1816229.shtml

Supercomputing Application Areas

The following graph shows the number of systems being used for the major application areas in china. From this it is clear that energy, gaming, scientific computing and meteorology are the primary application areas utilizing Chinese HPC systems.

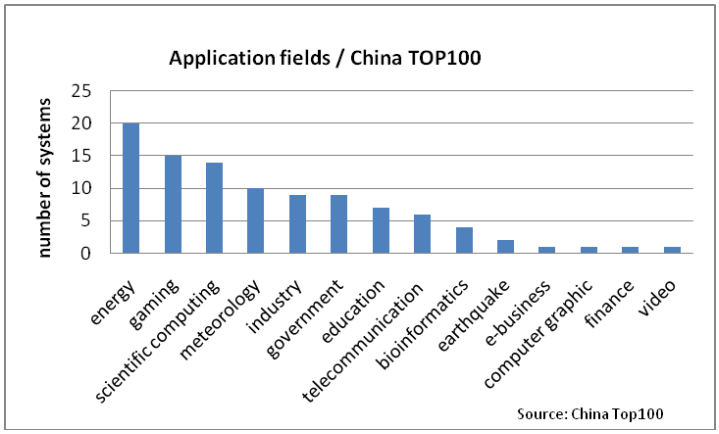

Figure 5: Application fields related to the Top 100 systems in China

The Chinese government has also assigned to its leading supercomputers centres specific missions. For example, the National Supercomputing Centre's mission is to provide support for meteorology, finance and other public information industries. The centre is also expected to provide innovation solutions for industry, provide support for advanced computing software and facilitate technology transfer and international cooperation in South China regions.[60]

Finally, the Dawning Nebulae, a hybrid GPU system, is planned to be used extensively by major Chinese Oil and Gas companies such as PetroChina (one of the richest companies in the world) and Sinopec, with real hybrid production codes.

The number of HPC service users in China is however relatively low. According to a recent report published by ZDnet China[61], most HPC services in China have less than 50 users and only 7.7% of all the services have a user group larger than 500 users.

Future near term Petascale Systems

The Chinese government is committed to developing the HPC ecosystem. The 12th five-year plan covering the period 2011-2015 is expected to be unveiled in the fall 2010 and will serve as a blueprint for the future social and economic reforms for the next few years. According to Dr. Kai Lu (School of Computer Science, National University of defense Technology), who gave a presentation at ISC'10 on the HPC market and research in China[62], the new plan is likely to include programs to develop not only HPC systems but also a service environment and applications. Dr. Kai Lu also highlighted that new technology will be adopted in domestic HPC systems (for example. the F-1000 processor and custom interaction technology for the TH system and the Loongson-3B processor for the Dawning system), and that the first 10 Petaflop machine could appear in China in 2013.

According to the Chinese Ministry of Science & Technology (MOST), a new major national Science and Technology project has been established to develop core electronic components, high-end

[60] http://english.peopledaily.com.cn/90001/90776/90881/6815913.html
[61] http://server.zdnet.com.cn/server/2009/0910/1461797.shtml
[62] Dr. Kai Lu, The *Developments of HPC Market & Research in China*, presentation at the International Supercomputing Conference (ISC'10), Hamburg, 1st June 2010.

generic chips, and generic computer programs. According to Prof. WEI Shaojun, the deputy chief designer of the project, quoted by MOST, China will build a proprietary Petaflop supercomputer, or Milky Way II, on the FT-1500CPU in 2011. The project was established to develop two components (core electronic components and generic software), and one chip (high-end generic chips). Researchers are currently working on a multi-core CPU that can be used for high performance computers. In the area of generic computer programs, researchers will focus on, for example, software at the operating system level, for databases and for middleware. China is striving to establish a proprietary research and development system with high-end generic chips and generic computer programs, to be realized in 2020.[63]

3.2.4 South Korea

Population: 50 M
GDP: 800 $M[64]
Yearly average supercomputer spending: 66 $M[65]
Supercomputer spending / GDP: 0,0083%[66]

HPC policy and funding programs

South Korea's R&D is centralized and coordinated by the National Science and Technology Council (NSTC), together with the Korean Ministry of Education, Science and Technology (MEST). In 2008, the Korean government refined its current Science and Technology Basic Plan which sets policy direction for science and technology for the period 2008-2012, listing 577 initiatives. The government has also drafted a mid to long term plan called the "Vision 2025: Development of Science and Technology" which sets the longer term main goals for science and technology[67].

In the IT field, a number of national initiatives aimed at developing this sector have been launched, such as the National e-Science Project (2005-2012) which is funded by the Ministry of Science and Technology and aims to build a next generation research environment, allowing researchers to share resources such as high-end research and large scale data facilities through an e-infrastructure. Four scientific areas are targeted: Life Sciences, Physics, Engineering, and Geosciences[68].

Supercomputing environment

South Korea has a long tradition of investment in HPC systems. In 2005, 14 systems hosted by universities, research institutes and government agencies were listed in the TOP500 list, placing Korea forth after the USA, Germany and the UK. Since then South Korea's position has declined – with only one system in 2010 – however the country has continued to invest in HPC systems. The Korean Institute of Science and technology Information (KISTI)[69] hosts a Sun system with a peak performance of 307.44 TFLOPS (274.80 TFLOPS Linpack performance) and which ranked number 15 in the TOP500 list in June 2010.

KISTI is currently the leading supercomputing center in South Korea. KISTI's Supercomputing Center is funded by the South Korean government and provides access to its supercomputing services, which operate on the highly advanced infrastructure of KREONET, to Korean university researchers, institutes, industries, and government organisations. In addition to all these services, the center also provides practical applied technology.

[63] http://www.most.gov.cn/eng/newsletters/2010/201003/t20100310_76233.htm
[64] Source: IDC report 2010
[65] Over the last five years. Source: IDC report 2010
[66] Source: IDC report 2010
[67] http://cordis.europa.eu/erawatch/index.cfm?fuseaction=ri.content&topicID=13&parentID=12&countryCode=KR
[68] http://www.ksc.re.kr/eng/project/project2.htm
[69] http://www.ksc.re.kr/eng/

Based on the National e-Science Project, *K e-Science* is one of the KISTI's major projects. It aims to build a leading research environment with advanced research organizations, to develop major technologies, and to build a global research infrastructure. By combining this program and the National e-Science Project, which aims to develop major technologies for e-Science and to build an e-Science testbed, the goal is to provide a world-class e-Science environment in the near future.

Upgrade plans

KISTI is at the early stages of developing its next supercomputing system. The goal is to install a system with a peak performance of 30 PFLOPS in 2013. If this were to transpire, South Korea would rank highly on the world stage. A key target area for this system is Nanotechnology.

The HPC community in Korea has been working to establish a law similar to the US High Performance Computing Act of 1991 in order to renew existing infrastructure and develop technology. Congressman Doo-Un Chung recently submitted a motion to create such an act which was supported by 17 other congressmen and which is being reviewed by the Committee on Education, Science and Technology.

Key component of the law include: to establish a national supercomputing committee; to formulate a comprehensive national plan for the entire supercomputing ecosystem (including infrastructure, application, education, and research) and to establish a national supercomputing center.

3.2.5 India

Population: 1.130 M
GDP: 1.243 $M[70]
Yearly average supercomputer spending: 19.6 $M[71]
Supercomputer spending / GDP: 0,0016%[72]

HPC policy and funding programmes

In India, responsibility for science and technology and R&D policy and funding is shared between central government and regional states. An important share of the funding and initiatives also comes from the private sector.

Policy setting for the R&D agenda and the priorities for science and technology areas of research are initiated through their Five Year Plan. The current XIth Five Year Plan (2007-2012) contains a section on Information Technologies that emphasizes the need to build on the existing strength of India in the software sector and to help develop from services to product development. The plant also calls for a strengthening of their IT hardware manufacturing base so that India can emerge as a regional hub for hardware manufacturing and export[73].

Another important document is the Report of the Steering Committee on Science and Technology for the XIth Plan issued by the Planning Commission and which sets the policy agenda for S&T and R&D for the period 2007-2012. In this document, the Planning Commission acknowledged that the Xth plan has enabled Indian computing power to go up from Gigaflops to Teraflops, and defined Petaflop computing as the next frontier and strategic are for India. The report also supported the establishment

[70] Source: IDC report 2010
[71] Over the last five years. Source: IDC report 2010
[72] Source: IDC report 2010
[73] http://planningcommission.nic.in/plans/planrel/fiveyr/11th/11_v3/11th_vol3.pdf

of an Indian Grid for Science & Technology (IGST) in order to develop state of the art IT technologies, computing and advanced computing modeling.[74]

Supercomputing environment

In November 2007, India made a remarkable entry into the Top10 list of most powerful supercomputers with a 172 TFLOPS system – called "eka", Sanskrit name for "one" – and which was developed by Tata's Computational Research Laboratories (CLR). This ranked forth in the TOP500 list. This entry reflected as much the dynamism of the Indian private sector in the development of HPC – "eka" is the first supercomputer developed by a private sector corporation and the fastest supercomputer available for commercial usage – as the relative absence of government funded large initiatives in India, although the initial plan partly originated in a national initiative supported by the Report of the Steering Committee[75].

Over the past few years, India has avoided funding large centers, funding a range of small to midrange HPC centers in various parts of the country[76]. The major government funded supercomputing centre in India is the Center for Development of Advance Computing (C-DAC). This was started in the 1980s to develop indigenous supercomputer capabilities. Other major national labs and research institutions include the Bhabha Atomic Research Center (BARC), the Advanced Numerical Research & Analysis Group (ANURAG) and the National Aerospace Laboratory (NAL).

In June 2010, India had five systems listed in the TOP500 list. The leading machine was the Tata 172 TFLOPS system which has dropped to a ranking of #33. Other systems – all under the 100 TFLOPS mark – are owned by the Indian Institute of Tropical Meteorology (70.39 TFLOPS peak, # 94), the Center for Development of Advanced Computing (54.01 TFLOPS peak, #182), IIT Kanpur (34,50 TFLOPS peak, #369), and an IT Services provider (53,09 TFLOPS peak, #389). At present, there are around 20 institutions that host Teraflop performance supercomputers[77]. These computers are primarily used in weather forecasting and climate modeling. Other important areas include computational physics, rendering, animation and computational fluid dynamics.

C-DAC is currently developing a Petaflop system, due for delivery in 2012[78]. This system is based on their PARAM series of systems, which are designed and built by C-DAC themselves. At a cost of $125M the system is expected to be used for weather forecasting, drug design, molecular modeling and remote sensing.

One of the strengths of India may reside in its software sector. One of the key elements of the growth of the software industry in India has been the Software Technology Parks of India (STPI) which are a result of public-private partnerships. Currently there are over 40 Software Technology Parks which produce 70% of software exports from India[79]. Some scientists believe that India may suffer from a lack of skilled manpower in the future to develop application software for supercomputer systems.

[74] http://planningcommission.nic.in/aboutus/committee/strgrp11/str11_sandt.pdf
[75] http://planningcommission.nic.in/aboutus/committee/strgrp11/str11_sandt.pdf, p.115
[76] IDC report 2010
[77] For more information: http://topsupercomputers-india.iisc.ernet.in/jsps/june2010/index.html
[78] http://www.cdac.in/html/press/3q09/spot687.aspx
[79] http://cordis.europa.eu/erawatch/index.cfm?fuseaction=ri.content&topicID=4&countryCode=IN

3.3 HPC initiatives in Europe

3.3.1 General overview

According to a recent study published by the market research firm International Data Corporation (IDC) and commissioned by the Information, Society, and Media Directorate General of the European Commission[80], Europe has under-invested in HPC compared to other nations, both in terms of resources and in research. However the same study also noted that "European stakeholders from industry, research, and academia believe that Europe has a chance to jump back to the forefront of development for the next generation of HPC-based research, and for the applications and other software technologies required for the transition to petascale and exascale computing".

Unlike the US and Asia, Europe does not have a large HPC manufacturing industry (with the exception of BULL) and has not sought to develop "domestic technologies" such as in Asia. Paradoxically, this can be viewed as an asset resulting in less dependence on national hardware vendors and allowing Europe to focus on other areas where it already excels (e.g. software development) and which might be central for future leadership in the global HPC market[81].

Europe is also in a unique position regarding its HPC policies and funding programmes. As with Asia, most countries have their own HPC strategy supported by national funding. Unlike Asia, these countries can also benefit from a strong overarching body (the European Commission) which coordinates EU wide strategies and funds leading-edge HPC initiatives, including those with neighbouring countries such as Russia.

A number of European countries already rank highly in the HPC market with systems competing with the most advanced in the US and Asia.

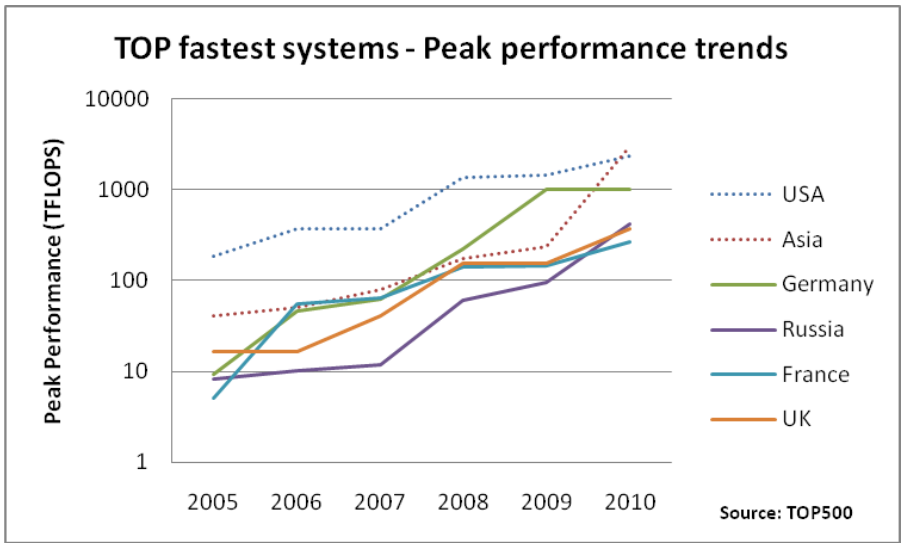

Figure 5: #1 systems (Peak performance) of the top 3 EU countries (Germany, France, and the UK) and Russia, compared with Asia and the US from 2005 to 2010 (Source TOP500 June lists 2005-2010)

[80] The development of a Supercomputing Strategy in Europe" (SMART 2009/005, Contract No 2009/S99-142914

[81] IDC report 2010

In 2009, Germany was the second country in the world to install a Petaflop supercomputer and hosts 24 systems in the TOP500 list (June 2010). Seven of these systems are above the 100 TFLOPS peak performance threshold. In France, the CEA recently announced the setup of the TERA 100 supercomputer, the first Petaflop machine (theoretical performance of 1.25 PFLOPS) to be fully designed and built in Europe by BULL. In addition GENCI will install a BULL Petaflop machine by the end of 2010, beginning of 2011, accessible for European researchers though the PRACE Research Infrastructure. The UK holds the largest share and number of systems in the top500, after the US. Other countries such as Spain, Italy, the Netherlands, Switzerland and Finland are also major worldwide players and are involved in ambitious HPC programmes, both at the EU level and nationally. Russia is another key actor with Petaflop plans and with whom the EU has recently launched joint HPC initiatives.

In this section, we review the EU wide and national initiatives related to HPC development towards Exascale and which have the potential to give Europe a leadership position in the road towards Exascale.

3.3.2 European Union

HPC policy and funding programs

In Europe, the largest HPC or HPC-related projects are funded under the FP7 "Capacities" program aimed at optimizing the use and development of research infrastructures. This runs until the end of 2013. The commitment of the European Union to support the development of research in HPC and to pursue HPC leadership has recently been demonstrated through several successful initiatives funded under this programme including PRACE, the current flagship project. On the software side, the programme funds key projects such as TEXT and it is just about to launch a call on exascale software and systems under ICT work programme 2011-2012.

Partnership for Advanced Computing in Europe - PRACE

The Partnership for Advanced Computing in Europe (PRACE) is a pan-European consortium comprising 20 countries and managed as a single European entity. It represents the top of the European HPC ecosystem, and provides Europe with world-class HPC systems.

PRACE's target is to maintain a pan-European HPC service consisting of up to six world-class leadership systems (Tier-0) integrated into the European HPC ecosystem. In the medium term, each system will provide several Petaflops of computing power and in the longer term (2019) the aim is for Exaflop computing power.

The PRACE Research Infrastructure is already in operation. The first production system, a 1 PFLOPS IBM Blue Gene/P system (JUGENE) was installed at Forschungszentrum Jülich (FZJ), Germany. As of June 2010, JUGENE ranked #5 in the TOP500 list.

Several other world-class systems will follow and are expected to be installed in France (in 2011), Spain, Italy, and the Netherlands. PRACE has to date secured 400 million Euros of national funding for Tier-0 systems and associated services for the next five years. These systems will be partly funded by the EU's seventh Framework Program through the First Implementation Phase project (PRACE-1IP) which started on 1 July 2010 and partly by national funding bodies.

Project name	**Partnership for Advanced Computing in Europe (PRACE)**
Short description/target	Build a world-class high-performance computing capability and strengthen Europe's scientific and industrial competitiveness. On the longer term (2019) Exaflop computing power will be targeted by PRACE
Partners involved	20 partners
Url	http://www.prace-project.eu

Funding source	PRACE members, European Commission
Schedule	PRACE Preparatory Phase: January 2008 – June 2010 PRACE First Implementation Phase: July 2010- June 2012
Budget	PRACE 1IP: 28 M€ (20 M€ EC contribution) Additional 400 M€ from hosting partners for five years

In 2007, a consortium for the "Promotion of Supercomputing Partnerships for European Competitiveness and Technology" (PROSPECT) was launched by FZJ, LRZ and BSC, aside of the PRACE project[82]. Under the PROSPECT umbrella the Parties – which now include several other supercomputing centers and universities – intend to create projects concerning the main objective to develop hardware and application and system software components as well as related operating and programming concepts for Peta- and multi-Petaflop systems as a contribution to the PRACE-project.

Distributed European Infrastructure for Supercomputing Applications (DEISA)

DEISA provides a pan-European HPC infrastructure for supercomputing applications. Consisting of eleven national supercomputing centres connected by a high performance network, DEISA provides a production quality environment for EU scientists to exploit. Funded under the Framework 7 program (and previously under Framework 6) DEISA is looking to demonstrate the possibility of a persistent European wide HPC infrastructure and environment.

The DEISA Extreme Computing Initiative (DECI) solicits grand challenge projects to utilise the DEISA infrastructure from EU scientists in all scientific areas. The aim is enhance DEISA's impact on science and technology.

Project name	Distributed European Infrastructure for Supercomputing Applications **(DEISA)**
Short description/target	Provide a pan-european HPC infrastructure for Supercomputing applications.
Partners involved	11 partners
Url	http://www.deisa.eu/
Funding source	European Commission, DEISA partners
Schedule	2002 – 2011 (DEISA2 from May 2008 – May 2010)
Budget	DEISA2: 18.7M€

European Exascaqle Software Initiative (EESI)

The goal of the European Exascale Software Initiative (EESI) is to build a European roadmap along with a set of recommendations to address the challenge of performing scientific computing on the new generation of computers composed of millions of heterogeneous cores which will provide multi Petaflop performances in 2010 and Exaflop performances by 2020. These hardware capabilities lead to outstanding technological breakthrough possibilities in computations and simulations which will be reached only if an international cooperation work program is set up. This will be achieved by organising a set of conferences and setting up work groups involving a very large number of HPC European actors, including both scientific software developers and users. They will investigate where Europe stands in the overall international HPC landscape, what are its strengths and weaknesses,

[82] http://www.prospect-hpc.eu/

what are the priority actions, and what type of cooperation structure should be implemented between Europe and the international community. EESI will also identify the on-going competitiveness in and for Europe induced by the use of Peta/Exascale software. It will investigate and propose education and training programs to be set up for the next generation of computational scientists

Project name	**European Exascale Software Initiative (EESI)**
Short description/target	Build a European Software Exascale Roadmap
Partners involved	8 contractual partners and 17 associate partners
Url	http://www.eesi-project.eu/.
Funding source	European Commission, EESI partners
Schedule	2010 – 2011
Budget	0,64 M€

Progress towards Exascale Applications (TEXT)

The European Commission is committed to Exascale software development. One of the key initiatives recently launched in this area, in addition to the European Exascale Software Initiative, is the Towards EXascale ApplicaTions (TEXT) project.

TEXT is an EC funded research project into new parallel programming models, in particular StarSs, from BSC (Barcelona Supercomputing Center), which aims to scale real-world applications to an Exascale performance level. The project includes researchers and institutions from Spain, Germany, France, Greece, Switzerland and the UK. Together these countries will undertake a series of tests, using petascale machines in various HPC centers across the EU and a challenging group of real-world applications, to determine if the MPI + StarSs model shows promise in supporting application development for Exascale. The project will also advance the status of tools and development environments for the hybrid MPI/SMPSs model and develop locality and load balance optimizations. Finally, the project will try to influence the evolution of future standards, in particular OpenMP.[83]

Project name	**Towards EXascale ApplicaTions (TEXT)**
Short description/target	TEXT is an EC funded research project into new parallel programming models, in particular StarSs, aiming at scaling real-world applications to Exascale performance level.
Organisations involved	Barcelona Supercomputing Center (BSC) – Spain (coordinator) High Performance Computing Center Stuttgart (HLRS) - Germany Forschungszentrum Julich (FZJ) – Germany EPCC at the University of Edinburgh – UK Foundation for research and Technology (FORTH) – Greece University of Manchester – UK Université de Pau – France Universitat Jaume I – Spain IBM Zurich (Switzerland)
Contact person	Jesús Labartata (BSC), Technical manager of the project. (jesus.labarta@bsc.es)
Funding source	EC and partners
Project duration	24 months, starting June 2010
Budget	3,5 M€ with an EC contribution of 2,47 M€ (call infra 2010-1-2-2)

EC call on Exascale computing, software and simulation

[83] http://theexascalereport.com/content/2010/eu-searches-new-super-model-text-project

"*A special objective on exascale software and systems with a significant international cooperation dimension, a long term perspective and the engagement of industry is foreseen under ICT Work Programme 2011-2012". This is the first objective in FP7 dedicated specifically to exascale computing.*

"The target is to develop a small number of advanced computing platforms (relying on vendors' proprietary hardware or on COTS) with potential for extreme performance (100 PFLOPS in 2014 with potential for exascale by 2020), as well as optimised application codes driven by the computational needs of science and engineering and of today's grand challenges such as climate change, energy, industrial design and manufacturing, systems biology, etc."

Applications for this call are due by January 2011 and will consist of groups of leading HPC centres, technology suppliers (industrial or academic) and industrial and academic centres. Interestingly, proposals should involve a rough split of 40/60 applications / systems development and focus on major challenges such as handling the expected extreme parallelism required by these systems.

The call is expected to place Europe at the forefront of HPC software development; strengthen European industry supplying and operating HPC systems; demonstrate European excellence in exascale level simulation and re-enforce international cooperation.

Project name /call	**Exascale computing, software and simulation (FP7-ICT-2011-7 Objective ICT-2011.9.13)**
Short description/target	The objective is to develop a small number of advanced computing platforms with the potential for extreme performance (100 PFLOPS in 2014 with potential for exascale by 2020), as well as optimised application codes driven by the computational needs of science and engineering..
Url	ftp://ftp.cordis.europa.eu/pub/fp7/ict/docs/ict-wp-2011-12_en.pdf
Funding source	European Commission
Budget	25 M€ (IP: 24M€ / CSA: 1M€)
Schedule	Projects funded for 2-3 years, starting in 2011-12. Deadline for applications: 18 January 2011

EC call on Computing Systems

Another ICT call published by the European Commission will fund projects to develop multi-core architectures. "*The objective is to achieve breakthroughs in the transition to multi-core architectures across the whole computing spectrum: embedded computing, general-purpose computing (PC/servers) and high-performance computing (HPC). This transition affects the underlying hardware, the system software (compilers, tools, OS, etc) and the programming paradigms*."

Research topics for this call will include Parallel and Concurrent Computing (multicore, multichip, parallel/concurrent software and tools, etc.), Virtualisation (heterogeneous multicore systems), Customisation (reconfigurable architectures, multicore on single chip, tool-chains, system modeling and simulation), Architecture and Technology (3D tacking, alternative computation models), and International Collaboration.

This call is expected to provide impact by: improving the programmability of future multicore systems; providing efficient virtualization; enhancing accelerated system development and production; re-enforce European excellence in multi-core architectures, software and tools and strengthen European leadership in cross-cutting technologies.

Project name /call	**Computing Systems (FP7-ICT-2011-7 Objective ICT-2011.3.4)**
Short description/target	"The objective is to achieve breakthroughs in the transition to multi-core architectures across the whole computing spectrum: embedded computing, general-purpose computing (PC/servers) and high-performance computing (HPC)."
Url	ftp://ftp.cordis.europa.eu/pub/fp7/ict/docs/ict-wp-2011-12_en.pdf

Funding source	European Commission
Budget	45 M€ (STERP: 40M€ / NoE: 4M€ / CSA: 1 M€)
Schedule	Deadline for applications: 18 January 2011

EU-Russia collaboration

Collaboration between the EU and Russia has recently been strengthened by the launch of the first joint EU-Russia ICT call (Objective ICT-2009.10.2 EU-Russia Research and Development cooperation) focusing on programming models and runtime support, performance analysis tools for HPC, and optimisation, scalability and porting of codes. Russia has leading academic expertise in a variety of scientific fields that use HPC, while the EU has leading academic expertise in parallel programming and is home to a dynamic ecosystem of innovative HPC tool companies. The call will include 3 stages of 2 years supported by a total of 6 M€ (4 M€ EU + 2 M€ Russia).

Russia has recently increased its spending on HPC and has a desire to become a major player in the HPC market. In March 2010, Russian Prime Minister Vladimir Putin allocated RUB 1.1 billion (28 M€) to develop supercomputer technologies in Russia, according to the Azeri News Agency APA[84]. President Medvedev has also called for a greater number of supercomputers and for a grid infrastructure in Russia.[85] Today, the Lomonosov supercomputer located in Moscow State University (MSU) – a 414 TFLOPS peak performance system with Linpack mark of 350 TFLOPS – is the fastest supercomputer in Russia and the second in Europe, ranking 13th in the TOP500 list of June 2010. The Lomonosov supercomputer will undergo a performance upgrade with the goal of obtaining a Petaflop level system in 2010. According to he MSU President, quoted in an HPCwire report, the anticipated cost of increasing the computer's performance to reach Petaflop level is around RUB 900 million (23 M€)[86]

The Lomonosov supercomputer was manufactured by the Russian HPC vendor T-Platforms, a rapidly growing private company established in 2002 to provide supercomputing systems, software and services in Russia. T-Platform has now installed almost 200 supercomputer systems and is currently expanding its activities into Europe and Asia. Russia is currently developing a next-generation HPC operating system, Clustrx, which is designed to enable the evolution from Petascale to Exascale[87]. In May 2010, T-Platforms was chosen to manage a $6M programme funded by RUSNANO, the Russian Corporation of Nanotechnologies[88], to enable commercial organizations in Russia exploit HPC resources and create a commercial market place for supercomputer simulations. Industrial applications will be selected from certain targeted industries including shipbuilding, aerospace, automotive, oil and gas, chemistry, pharmaceuticals, energy and construction.[89]

In the past few years, a number of national programmes and partnerships have been launched to develop HPC technology and systems in Russia, such as the supercomputer program "SKIF GRID", sponsored by the Union State of Russia and Belarus, in collaboration with T-Platforms, the Russian Academy of Sciences and Intel Corporation. This programme was initiated in April 2007 for three years by the government with an aim of fostering and supporting technology development in the HPC and grid field. Several SKIF Teraflop computers have been developed under this program (including a 60 TFLOPS system - SKIF MSU – deployed at Moscow State University). The programme also fostered new partnerships in Russia. As an example, the SKIF-Aurora project, involved a consortium composed of the Russian Academy of Sciences and the HPC system integrator RCS SKIF and aimed to deliver Petaflop and Exaflop-scale supercomputer solutions. The first SKIF-Aurora installation (24 TFLOPS) was installed in 2009 in South Ural State University (SUSU).

[84] http://en.apa.az/news.php?id=117216

[85] http://enterthegrid.com/primeur/09/articles/weekly/AE-PR-09-09-53.html

[86] http://www.hpcwire.com/features/Moscow-State-University-Supercomputer-Has-Petaflop-Aspirations-88264637.html

[87] http://www.t-platforms.ru/images/pdf/Clustrx_OS_Clustrx_Watch_20-05-10.pdf

[88] http://rusnano.com

[89] http://www.t-platforms.ru/en/about/allnews/236-rosnanonews.html

3.3.3 National initiatives in EU Member States

In this section we describe the current national funding bodies and programs for the major EU countries involved in HPC. Future Petascale platforms are highlighted and relevant Exascale initiatives are described. Several European countries have launched ambitious government HPC hardware and software programmes. Some of them have also joined forces to develop hardware and application and system software components under the PROSPECT consortium[90].
In terms of Exascale initiatives, the most significant development are from public-private partnerships and collaboration between research centres and industry. Some of the most recent and interesting initiatives are reviewed in sub-section 2.

HPC national and regional programmes in Europe

A recent IDC report on "The development of a Supercomputing Strategy in Europe" (SMART 2009/005, Contract No 2009/S99-142914) summarised the nature of HPC funding in various counties and we acknowledge this source as a starting point for much of the detail in the following section.

The United Kingdom

The Engineering and Physical Sciences Research Council (EPSRC)[91] is the managing agent for the UK Research Councils HPC programme. This is run on behalf of the Office of Science and Technology (OST) with the High Performance Computing Strategy Committee (HSC) advising the Research Councils and OST on HPC strategy. EPSRC is currently the managing agent for the UK National Service, HECToR[92], based at the University of Edinburgh, ranked number 16 in the top500. An upgrade in 2011 will take this system to between 600 and 900 TFLOPS. A strategic framework for High End Computing (HEC) was produced in 2007 by the High End Computing Strategic Framework Working Group at the request of the High End Computing Strategy committee. This framework helps guide HPC developments in the UK including procurement, access and promotion. EPSRC currently have a call for HPC Software Development projects (August 2010), to enable HPC science and engineering, with up to £3.5M available.

Three organisations (EPCC, STFC and NAG) provide national service support for high performance computing. EPCC at the University of Edinburgh[93] operates the UK National Service, HECToR, a 134?? euro project with a 374 TFLOPS Cray XT6, soon to be XE6, system. The centre has a project based consultancy service, with over 250 companies having benefited from their products and services. STFC[94] is currently building on its Computational Science and Engineering department's HPC expertise to develop the Hartree Centre, a new computational sciences institute for the UK. This exciting project will bring together academic, government and industry communities to focus on multi-disciplinary, multi-scale, efficient and effective computation. NAG[95] provides the Computational Science and Engineering (CSE) Support Service for HECToR, providing dedicated optimisation and scaling support and HPC training for users of the service. They also provide CSE services and HPC technology consultancy to businesses. Finally, the Oxford e-Research Centre[96] has recently coordinated an activity to develop a High Performance Computing / Numerical Algorithms roadmap for the UK, to gather and understand the common algorithmic requirements from many different fields and to suggest future development scopes and targets.

[90] http://www.prospect-hpc.eu/
[91] http://www.epsrc.ac.uk/funding/facilities/hpc/
[92] http://www.hector.ac.uk
[93] http://www.epcc.ed.ac.uk
[94] http://www.cse.scitech.ac.uk/
[95] http://www.nag.co.uk/
[96] http://www.oerc.ox.ac.uk/

The Atomic Weapons Establishment (AWE)[97], the Met Office[98] (the UK's National Weather Service) and the European Centre for Medium-Range Weather Forecasts (ECMWF)[99] all host significant HPC resources within the UK.

France

HPC policy for research is determined by three sources in France: a strategic committee called "le Comité Stratégique du Calcul Intensif" (CSCI) which was created in 2005; the Grand Equipment National de Calcul Intensif (GENCI), created in 2007, which allocates HPC resources to French researchers, coordinates and funds the facilities of the 3 national civilian research computing centres and finally the Agence Nationale de la Recherche (ANR) which has funded more than 66 M€ of research in HPC since 2005 through dedicated programs (like Cosinus).

The 3 national HPC centres (CNRS/IDRIS, CEA/CCRT and CINES) operate the national resources for research applications. As of 2003, CEA/CCRT has also opened its resources to industrial users through partnerships that give them access to the same kind of computing resources and support for applications available to the research community.

In order to facilitate HPC usage by French SMEs (Small and Medium sized Enterprises) GENCI, INRIA and OSEO launched a joint initiative called "HPC pour les PME" in September 2010. This initiative provides an integrated offer based on training, expertise, access to HPC facilities and funding to assist SMEs to assess the potential of HPC in their daily business in order to increase their competitiveness. CEA is building a new facility called TGCC (Très Grand Centre de Calcul), their first hosted machine will be the French tier-0 PRACE machine, a Petaflop system funded by GENCI, in 2011.

In addition to HPC research carried out by these centres, INRIA (national computer science research organisation) has an extensive HPC activity with several teams involved. As an example in June 2009, the University of Illinois at Urbana-Champaign and INRIA formed the Joint Laboratory for Petascale Computing. The Joint Laboratory is based at Illinois and includes researchers from INRIA, Illinois' Center for Extreme-Scale Computation and the National Center for Supercomputing Applications. It focuses on software challenges found in complex high-performance computers.

HPC is also represented in multiple competitiveness clusters (like System@tic in the Paris area) through collaborative R&D projects between academia and industry fostering technological transfer. Three French industrial companies have strong HPC policies of their own and large computing facilities, with Petascale prospects: Airbus (aeronautics), EDF (energy) and Total (oil & gas).

Ter@tec is a French non-profit association for promoting HPC which links industrial companies, users and providers, and major R&D centers, public and private. Ter@tec's main objective is the creation of a wide ecosystem dedicated to simulation and HPC. To achieve this goal, it is acting along three axes: participate in the emergence and setting up of joint collaborative R&D projects, facilitate the access to state-of-the-art processing facilities among the most powerful in the world and contribute to the development of the requisite skills and competences through training, seminars, conferences and lectures.

[97] http://www.awe.co.uk/

[98] http://www.metoffice.gov.uk/

[99] http://www.ecmwf.int/

Ter@tec is planning to install the first technopole in Europe dedicated to Simulation and High Performance Computing. The Ter@tec Campus (13 000 m2 of offices and labs), close to the TGCC (Very Large Computing Center), will host more than 1000 people, including the major players in the HPC industry (manufacturers, publishers, suppliers of service), some collaborative laboratory Industry & Research (Exatec HPC Lab involving Intel Corp, Bull-CEA Extreme Computing Lab, …), an incubator and nursery hotel companies, some Services Platforms (hardware, software and services) and a HPC European Training Institute.

Created in 1994, the ORganisation Associative du Parallélisme (ORAP, www.irisa.fr/ORAP/) association has actively promoted HPC at the national level by organizing a conference every six months and by editing a newsletter to spread information related to HPC among the French public and private organizations.

While collaborating with many of the aforementioned actors, CEA/DAM has its own HPC roadmap for defence applications, In May 2010, CEA/DAM announced the setup of its TERA 100 supercomputer, the first Petascale machine (theoretical performance of 1.25 PFLOPS) to be fully designed and built in Europe. The TERA 100 program is a close collaboration between BULL and the CEA in the design and development of new Extreme Computing technologies.

Germany

High Performance Computing policy in Germany is decided at both a federal and state level with advice from the Wissenschaftsrat (WR). A WR committee provides approval for significant HPC investments while a German Research Council (DFG) committee addresses the provision of supercomputing resources to Universities and makes recommendations to the WR. In general, academic HPC systems are jointly funded by the federal government and by the state where the University is located. For example the BlueGene/P system at the Forschung Zentrum Jülich (FZJ) is funded by the federal state of North Rhine-Westphalia, the Federal Ministry of Education and Research (BMBF) and the Helmholtz Association.

The Gauss Centre for Supercomputing (GCS) was created in 2006 under an agreement between the BMBF and the regional ministries of North Rhine-Westphalia, Bavaria and Baden-Württemberg. This centre is coordinating the collaboration between FZJ, the Leibniz-Rechenzentrum (LRZ) and the Höchstleistungsrechenzentrum Stuttgart (HLRS). The supercomputing activities of GCS are funded with 400 M€ by the national project PetaGCS (2010 – 2016) which aims to establish 3 Petaflop supercomputers at its member sites. While the Gauss-Centre is primarily addressing the needs on the top end, regional and topical German HPC centres are teamed up by the Gauss Allianz to create the organizational framework and ecosystem necessary to coordinate the HPC related research activities of its members. Further goals of the Gauss-Allianz are the furtherance of computational sciences in order to increase the visibility to compete on an international level and to maintain the leading role of Germany in the European supercomputing activities.

In January 2010 the German federal government issued a call for “HPC Software for Scalable Computers” to improve the scalability of applications.

The Juelich Supercomputing Centre (FZJ) has recently signed two partnership agreements with IBM, Intel and ParTec to develop hardware and software for an exascale supercomputer (ECL), and key technologies, tools and methods to power multi Petaflop and Exaflop machines, focusing on the scalability and resilience of those systems (EIC) – see details in the following sub-section.

The Max-Planck Society, the Ludwig-Maximilians University Munich, the Technical University Munich, the Leibniz Supercomputing Centre (LRZ) of the Bavarian Academy of Sciences and Humanities and the Rechenzentrum Garching (RZG) of the Max-Planck Society created the Munich Computational Sciences Centre (MCSC). MCSC aims to foster joint research developments in the field of computational sciences, management of huge data repositories and HPC. The Munich Centre of Advanced Computing (MAC) is a research consortium which has been established at Technical University of Munich (TUM) to bundle research activities related to computational science and

engineering (CSE) as well as high-performance computing (HPC) - across disciplines, across departments, and across institutions. In MAC, seven of TUM's departments and other Munich research institutions (Ludwig-Maximilians-Universität, Max-Planck institutes, the Leibniz Supercomputing Centre of the Bavarian Academy of Sciences and Humanities) as well as TUM's international partners such as the King Abdullah University of Science and Technology (KAUST), join their forces to ensure the sustainable usage of current and future HPC architectures for the most relevant and most challenging CSE applications.

Fostering the use of modern supercomputers within scientifically and industrially oriented research communities to solve scientifically challenging and technically highly relevant problems has been the main focus of the Competence Network for Technical and Scientific High Performance Computing in Bavaria (KONWIHR).

Together with debis Systemhaus GmbH, debis-sfr and Porsche, HLRS has formed a company to provide access to supercomputers for research and industry.

Spain

The Ministry of Science and Innovation (MICINN), within the Secretariat of State for Research, governs Spanish national HPC policy and strategy. Some Regional Governments also have their own HPC policy and systems. Those systems are part of the Tier-2 level systems in Spain.

The Spanish National Facility in Supercomputing is the Barcelona Supercomputing Center (BSC), whose stakeholders are MICINN, Generalitat de Catalunya and the Technical University of Catalunya. The mission of the Barcelona Supercomputing Center - Centro Nacional de Supercomputación (BSC-CNS) is to facilitate scientific progress through research, development and management of information technologies. Its objectives are focused on scientific excellence in computer sciences and applied computing in life sciences and earth sciences, providing supercomputing support to Spanish e-Science, and creating wealth. It has, therefore, been recognised as a Unique Scientific and Technical Infrastructure. It represents Spain in PRACE and will be hosting a PRACE Tier0 system.

The BSC-CNS investigates systems hardware and software for the supercomputers of the future, and its staff includes researchers specializing in computer architecture, heterogeneous architecture, autonomous computing, grid computing, performance tools, programming models, network processors and storage systems. Other researchers are specialists in bio-IT, computational genomics, molecular modelling, computational proteomics, climate change modelling, meteorological modelling, air quality prediction and the atmospheric transport of mineral dust and volcanic ash. It also has experts in optimizing, scaling and customizing codes for efficient use in high-performance computers.

The BSC-CNS houses the MareNostrum supercomputer, one of the most advanced supercomputers in the world, which can currently carry out 94 billion operations per second thanks to its more than 10,000 processors. MareNostrum has been, since 2007, the central hub of the Spanish Supercomputing Network (RES, Red Española de Supercomputación). The Spanish Supercomputing Network (RES) is a Ministry of Science and Innovation initiative that meets the intensive calculation needs of the Spanish scientific community. It is an infrastructure of supercomputers distributed throughout Spain, which can be accessed on a transparent basis by the Spanish scientific community. This was unveiled in 2007 as a distributed supercomputing center and currently has eight nodes. Nowadays the Network is under expansion, incorporating new nodes and adapting the functional structure.

Joint collaboration R&D projects with private industry are undertaken at both the system level, via direct collaborations with industry leaders such as IBM and Microsoft, and at the application level, where a significant research effort is dedicated to providing supercomputing based models for solving engineering problems in industrial sectors such as aerospace, transport, energy, medicine, geology etc.

In 2008 a new project (*Supercomputación y e-Ciencia*) coordinated by BSC and with 5 M€ funding from the Ministry of Education and Science under the programme Consolider – Ingenio 2010 was launched with the aim of advancing supercomputing research through cooperation between researchers in Spain. The project which is going from 2008 to 2013 also aims The main aim of this

Consolider proposal is offering a national framework for research groups expert in supercomputing applications to collaborate together with expert hardware/software machine designers in order to design and use these machines efficiently in the near future.

Italy

HPC funding in Italy is generally supported by the Ministry for Education who provided 35 M€ in HPC funding for 2009.

CINECA is the Italian national HPC centre. It is based in Bologna and is operated by a consortium including the Ministry for Education, the Ministry of University and Research (MUIR), the Agency for New Technologies, Energy and the Environment (ENEA), the National Research Council (CNR) and the National Institute of Oceanography and Experimental Geophysics (OGS). CILEA in Milan, mainly supports Lombard universities while CASPUR in Rome, serves the universities of Rome and of South Italy.

The Euro-Mediterranean Centre for Climate Change (CMCC) is the national HPC Centre on Climate Change – it is involved in IESP and also chairs the WCES Working Group on Weather, Climate and Earth Sciences in EESI. The Centre is operated by a Research Consortium consisting of different Italian public and private research Institutions. The Centre has been established with initial funding by three Italian Ministries: Economy and Finance (MEF); Environment, Land and Sea (MAT); University and Research (MIUR). CMCC is involved in the development and in-depth examination of knowledge on climate variability and the study of causes and consequences through high resolution simulations. CMCC inaugurated its Supercomputing Centre located in the University of Salento in January 2009. It has two last-generation supercomputers making it one of the most important facilities in Europe for the realization of scenarios on the effects of climate change. The aggregated computational power available at the Centre is about 30 Tflops (theoretical peak performance) while the storage capacity is 1,5 PetaBytes. Such a supercomputing facility enables the CMCC to address challenging research issues and studies.

Finland

The main funder of HPC research in Finland is the Finnish Funding Agency for Technology and Innovation (TEKES) which funds academic and research organisations, generally supporting innovative projects. A flagship five year project called MASI aims to develop modelling and simulation methods and has a total budget of 26.3 M€. The Academy of Finland also funds HPC projects. The Finnovative Computing project (2010-2011), which is a sequel to the general code optimization project FinHPC, seeks to develop new programming models and investigates best practices for GPU accelerated parallel systems. Both projects are funded by Tekes. The SimITER project, funded by the Academy of Finland, focuses among other things on improving the performance of simulation codes for the Finnish computational fusion research community.

CSC – IT Center for Science – is the Finnish national HPC centre. It is responsible for providing national supercomputing and networking services and support and collaborates with industry on a wide variety of projects. In addition to maintaining and developing the Finnish HPC resources, CSC also has a strong emphasis on HPC programming, and runs a general code optimization service for its customers.

CSC operates a CRAY system featuring a theoretical peak performance of 102 TFLOPS, and an HP cluster system of 34 TFLOPS. In September 2010, CSC announced that it plans to build a new data center in Kajaani (Northern Finland), on the site of a former paper mill[100]. The data centre will be built jointly by CSC and UPM, and it will provide a state-of-the-art environment for supercomputers, data storage, and other demanding IT systems. The construction work of the data centre will begin during

[100] http://www.csc.fi/english/csc/news/news/data

the autumn and the centre will be completed in early 2012. This data centre is expected to be one of the most eco-efficient data centres, using local energy and water-cooling from the river. Around 30 M€ have been allocated to CSC to build this data centre, which also includes the purchasing of a new system.

Netherlands

In the Netherlands, the National Computing Facilities foundation (NCF) is responsible for academic supercomputing policy. NCF is also responsible for promoting and coordinating the use of the Netherlands' HPC and grid resources. NCF is an organisation under the NOW, the major funding agency for scientific research.

The NCF policy focuses on providing a long-term perspective for highly qualified scientific research by making sustainable and advanced computer facilities available to all researchers. Every six years, investment should be made in a top-of-the-line supercomputer, which should be upgraded mid way through the project. This long-term perspective is important for scientific research to benefit from developments in new methods and techniques. Developments in grid infrastructure will reinforce the need for top facilities such as supercomputers, large compute clusters, large data storage and middleware, as well as providing a match between these facilities and the local facilities. The starting point for the policy is both to stimulate research by providing such facilities as well as the establishment of a national e-Infrastructure making it possible to use these advanced systems.

The current national HPC system is the Huygens system, an IBM Power6-based system with a peak performance of more than 60 TFLOPS. The system consists of 32-core wide shared memory nodes. The system has an outstanding memory and I/O configuration, which has proved to be very valuable for the application set being run by researchers on this national capability system. In order to make sure that the right science is carried out on the right architecture, NCF (with some Dutch universities) has invested in a large compute cluster, which typically provides raw compute power, without the advanced features of the Huygens system.

In addition, the national grid project for e-science (called the BiG Grid project, of which NCF is the corresponding author, together with the national sub-atomic physics laboratory Nikhef and the bioinformatics community NBIC) invests in compute clusters, large storage facilities, network facilities to access the systems and large backup facilities. The project runs until 2011/2012, after which it is expected to be continued in a national research infrastructure environment.

A very important actor in the HPC arena is SARA, the Dutch national supercomputing center. SARA operates high-end HPC facilities, such as Huygens and the national compute cluster. SARA demonstrates expertise in all HPC areas, ranging from system administration to optimization and parallelization of user applications.

Another actor in the HPC arena is ASTRON, which runs the Lofar project, together with the University of Groningen. From an HPC perspective, a large IBM BlueGene/P system is used to perform the correlation of the received radio signals in the Lofar array of receivers in the Netherlands and neighbouring countries. Through NCF, users can obtain access to the BlueGene system as well.

NCF publishes on a regular basis for both the national HPC community and the international HPC and grid community. Well-known publications include those on recent supercomputers and on academic supercomputing in Europe.[101]

[101] "Overview of recent supercomputers", by Aad van der Steen, NCF, annual publication, and "Academic Supercomputing in Europe", by Rossend Llurba and Ana Bela sa Dias.

Public-private partnerships and collaboration

Over the past two years, public-private partnerships and collaboration targeting exascale and involving national research centers and international HPC companies have developed in several European countries. These collaborations have grown from a need for HPC centres to address their development needs and by major HPC companies wishing to develop their research network (Intel Labs, IBM collaboratory, Cray Research Initiative, etc.) and strengthen their position on the European market. In both cases, this brings additional funding to Europe for HPC research and contributes to the advance of European expertise in HPC hardware and software development. This section presents a summary of the most significant initiatives in this area.

Exascale Innovation Center (EIC) – Germany

On 23 March 2010, Forschungszentrum Jülich (FZG) and IBM signed a contract for a joint "Exascale Innovation Center" (EIC), located on its campus. EIC will develop hardware and software for an exascale supercomputer by the end of this decade. A prototype exascale system is anticipated by 2015. Five scientists from the IBM development laboratory in Böblingen and five scientists from Jülich will collaborate with a team of scientists at the IBM Watson Research Center in Yorktown Heights. The goal is to install an exascale-class system in Jülich by 2019[102].

Project name	**Exascale Innovation Center (EIC)**
Short description/target	Public-private collaboration to develop hardware and software required to approach exascale computing
Organisations involved	FZG, IBM
Schedule	2010-2019

ExaCluster Laboratory (ECL) – Germany

In May 2010, FZG also signed a multi-year agreement with Intel, and ParTec to create a new ExaCluster Laboratory (ECL) located on the campus of FZG. The second supercomputer based at FZG (JUROPA) is based on Intel's Nehalem processors and ParTec's ParaStation cluster middleware. ECL will initially employ about a dozen researchers and is expected to triple its staff over time. The lab is expected to develop key technologies, tools and methods to power multi Petaflop and Exaflop machines, focusing on the scalability and resilience of those systems[103]. Current challenges in systems management software for large heterogeneous supercomputer systems will be investigated. This will include research on open exascale runtime system software, software tools and simulation software[104]. The aim of the work on systems management software is to further improve the scalability of ParTec's ParaStation cluster middle-ware and – at the end – to create an OpenSource software stack capable to manage Exascale systems. ECL will investigate innovative technologies to overcome the yet unsolved problem of coupling accelerators more tightly to the high-speed interconnects of today's cluster computers.

Project name	**ExaCluster Laboratory (ECL)**
Short description/target	Public-private collaboration to develop key technologies, tools and methods to power multi Petaflop and Exaflop machines, focusing on the scalability and resilience of those systems
Organisations involved	FZG, Intel, ParTec
Schedule	Multi-year agreement

[102] http://www.fz-juelich.de/jsc/files/docs/newsletter/jscnews-184.pdf
[103] http://www.intel.com/pressroom/archive/releases/20100531comp.htm
[104] http://www.fz-juelich.de/jsc/docs/newsletter/2010/jscnews-186#ecl

EX@TEC – France

In November 2009, CEA, GENCI, the University of Versailles at Saint Quentin-en-Yvelines (UVSQ) and INTEL created a European laboratory for research on exascale computing, EX@TEC. EX@TEC will contribute to the development of hardware and software required for exascale computing. CEA, GENCI and UVSQ will share the funding together with Intel. This will represent an investment of several million Euros over a period of five years. This public private partnership combines the strong vision of French participants in scientific computing and the academic HPC software expertise knowledge of Intel. The activities of the laboratory include the exploration of the challenges related to the integration of multi-Petaflop systems, the optimisation of software performance and collaboration between the developers and users of large systems, in order to optimize applications and to anticipate upcoming hardware and software breakthroughs with Exascale. The energy, seismology, fluid dynamics and health sectors are involved in this collaboration. The EX@TEC lab kicked off with close to 20 people in early 2010.

Project name	**Ex@tec**
Short description/target	Public-private collaboration to develop hardware and software required to approach exascale computing
Organisations involved	CEA, GENCI, UVSQ, INTEL
Schedule	2010-2015

Flanders ExaScience Lab – Belgium

On 8 June 2010, Intel Corporation, IMEC and 5 Flemish universities officially opened the Flanders ExaScience Lab at the IMEC research facilities in Leuven, Belgium. The lab will develop software to run on Intel-based future exascale computer systems. The Lab will focus on power reduction and reliability, areas of critical importance for achieving Exascale supercomputer architectures. The Flanders ExaScience Lab will focuse on enabling scientific applications, such as the simulation and prediction of "space weather," or electromagnetic activity in the space surrounding the Earth's atmosphere. The Flanders ExaScience Lab has close to two dozen researchers and will add another dozen by 2012. The lab will be hosted at IMEC and is supported by the Flemish Government agency for innovation by science and technology (IWT).[105]

Project name	**Flanders ExaScience Lab**
Short description/target	Public-private collaboration to develop software to run on Intel-based future exascale computer systems.
Organisations involved	Intel Corporation, Imec, and five Flemmish Universities (University of Antwerp, Ghent University, Hasselt University, Katholieke Universiteit Leuven and Vrije Universiteit Brussel)
Schedule	2010-2015

Exascale Stream Computing Collaboratory – Ireland

IBM is establishing an exascale stream computing research ‘collaboratory’ in Dublin. The agreement will see IBM supercomputing and multidisciplinary experts work directly with University researchers from Trinity College Dublin (TCD), Tyndall National Institute in Cork, National University of Ireland Galway (NUIG), University College Cork (UCC) and IRCSET, the Irish Research Council for Science,

[105] http://www.intel.com/pressroom/archive/releases/2010/20100608corp.htm

Engineering and Technology to “develop computing architectures and technologies that overcome current limitations of dealing with the massive volumes of real-time data and analysis”. The exascale research in Dublin will focus on how Exascale systems can be applied to solving complex business problems. The research will include both technical and applications research.

Project name	**Exascale Stream Computing Collaboratory**
Short description/target	Public-private collaboration to develop computing architectures and technologies that overcome current limitations of dealing with the massive volumes of real-time data and analysis.
Organisations involved	IBM, Trinity College Dublin (TCD), Tyndall National Institute in Cork, National University of Ireland Galway (NUIG), University College Cork (UCC) and IRCSET, the Irish Research Council for Science, Engineering and Technology

Exascale Technology Centre – UK

In January 2010, Cray Inc and EPCC at the University of Edinburgh launched the Exascale Technology Centre, which joins Cray’s wider Exascale Research Initiative to will explore new ideas and new technologies to meet the challenge delivering an Exaflop within the next decade. Working closely with Cray research and development, staff from EPCC and Cray will be investigating new and novel programming languages for exascale; optimal algorithms and libraries and how to program accelerator based architectures effectively.

Project name	**Exascale Technology Centre**
Short description/target	Public-private collaboration to investigate the challenge of delivering an Exaflop within the next decade.
Organisations involved	Cray Inc and EPCC, The University of Edinburgh
Schedule	2010-2012

4. Conclusions

This report provides details of large-scale HPC initiatives in the US, Asia and Europe and examines trends and future developments in this area. Current Petascale initiatives are described and future projects looking at Exascale development considered. The goal is to understand Europe's current position in the international HPC landscape, understanding our strengths and weaknesses.

Following significant investment over the last few years, the US currently dominates the supercomputing sector, hosting two of the three Petascale systems[106]. In addition, significantly more are expected in the near future, including the HPCS prototypes. The US has also led much of the discussion around Exascale computing, investing in the International Exascale Software Project (IESP) and funding programmes such as the Ubiquitous High Performance Computing program and the co-design centre call. However, currently no multi-agency national Exascale initiative exists to drive the Exascale agenda forward.

China has invested heavily in supercomputing technology and now holds second position in the TOP500 list. It is likely that this funding will continue and China represents a serious contender to own the first Exascale system. The Chinese government has actively promoted independent innovation to avoid reliance on foreign technology. Although China's recent top-end systems are impressive, it is notable that the peak performance / sustained performance ratio is high, leading to questions over how applicable these systems are.

Japan and South Korea both expect to build Petaflop systems. Japan will have their first system by the end of 2010, and has a flagship project to develop a 10 Petaflop system by 2011-2012. South Korea has announced plans for a Petaflop system by 2013. Although India also plans to have a Petaflop system by 2013, funding has recently been directed at small to medium range HPC centres.

BULL is the only European HPC manufacturer. BULL will deliver two Petaflop systems in France respectively to CEA DAM, in 2010, and to GENCI in 2011. Up to now, Europe has not sought to develop it own "domestic technologies" such as in Asia. Europe has focussed on demonstrating expertise in software development and has recently seen a series of public-private partnerships between European research centres and international HPC producers to address the challenges associated with Exascale. Significant funding has been made available in Europe through the Framework 7 programme, both in Petascale (e.g. PRACE) and Exascale terms (e.g. the current Exascale computing, software and simulation call).

Most Exascale initiatives across the continents identify a similar set of global issues as scientific drivers for Exascale computing. These include climate change, sustainable energy sources and nanotechnology. The recent G8 Research Councils Initiative on Exascale applications for global issues is significant as it provides funding to enable collaboration between US and European partners and that it focuses on the applications of Exascale rather than hardware.

The US, Europe and China all have the potential to realise the first Exascale system. Europe distinguishes itself from other continents and countries by its lack of a home grown HPC manufacturing industry and its focus and expertise on software and support. While many countries have identified similar global issues as scientific drivers for Petascale and Exascale systems, the path to enabling these applications, at all software levels, is uncharted, providing Europe with the opportunity to drive this agenda.

[106] This figure is based on sustained, LINPACK, performance.

This investigation will be revisited and updated during the last four months of the project, mapping changes in the HPC landscape across the duration of the project.

