Towards A European Roadmap for Datacentres and their Economic and Environmental Impact

Report of a Workshop held in Brussels, November 2011

Executive Summary

Information is at the heart of every economy, developed or emerging. We use it for business, information services, telephony, social networking, image storage, science and a panoply of other applications. Today we hold equivalent of about 100 newspapers for every man, woman and child on the planet and the amount is growing rapidly.

This data-dependency has lead to the creation of large scale data-storage facilities, datacentres, to store, replicate and manage it. They each can consume vast amounts of power, some using the equivalent of many thousands of households. Not only do they consume energy and generate heat, but also have substantial water- and carbon footprints. As a result they have a large environmental and economic impact.

It is a global issue that Europe has begun to address through strategies including Europe 2020 and policies such as the Code of Conduct for Data Centres. In its stated aim to “Go Lean, Go Efficient and Go Green” it has the capability to lead the world.

This report sets out an integrated roadmap to help Europe address both the economic and environmental impact through a combination of Policy, Strategy, Economic, Business and Technology measures.

The result of a workshop “Towards A European Roadmap for Datacentres and their Economic and Environmental Impact” bringing together industry, technologists and the Commission, it identified the need for a potential EU programme to develop the EU ecosystem in this critical area through:

- **An integrated programme to improve resource efficiency** by at least 20% in a four to five-year timespan by attacking a series of “low-hanging fruit”;
- **A longer term technology-driven programme** which would aim to achieve up to 100x improvements in resource-efficiency over the next decade through new resource-efficient, low-cost systems
- **Actions targeted at datacentres efficiency**, not only at those of large corporations but also of SMEs, whose aggregated resource-usage has a very substantial environmental and economic impact;
- **Tools** to permit accurate measurement and monitoring of datacentre efficiency and regulatory frameworks to enable their users to monitor them.

Together these will enable an integrated approach to reducing the economic and environmental impact of datacentres.

CONTENTS

Introduction

We live in a society in which data has become central to our economic and social well-being. It seems incomprehensible to those growing up today that Society ever lived without the levels of access to information of all kinds that they take for granted. The coming years will see an exponential growth in the amount, types and sources of data that we store.

What began with Guttenberg has passed into the Internet Era as central to everything that we do, from how we conduct our daily lives, to how the European Union works. Data begets data; information is data in context; and information begets yet more information. The so-called “information explosion” of the last fifteen years has generated a major shift in the way that we use and store data and the quantity that we store.

The rapid increase in data storage has, in turn, led to a dramatic growth in the number of datacentres, facilities that primarily contain electronic equipment used for data processing, data storage, and communications networking. These vast repositories, often replicated many times for security or ease of access, have become at once a benefit and an ever-growing cost. Their consumption of electricity and the impact that they have on the environment is growing ever larger.

In November 2011 the European Commission held a workshop Brussels called “Towards A European Roadmap for Data Centres and their Economic and Environmental Impact” (“the Workshop”), convened to consider forward, ways in which Europe can minimise both the environmental and economic impact of such centres. This had as a starting point the recommendations of an earlier workshop held in Brussels in June 2011 to consider the technological issues[i]. The remit of the November workshop was broader than the June Meeting, considering not just technological issues, but also looking into drivers, and service and business models.

To provide a sense of scale: in 2007 datacentres in Western Europe consumed some 56 TWh of power per annum and according to some models will rise to about 105 TWh in 2020[ii, iii], other figures such as those provided by US Environmental Protection Agency Energy Star Programme[iv] suggest a figure for the US of some 300 TWh in 2011, of which as much as 60% may be consumed by volume data servers. More recent, informal estimates suggest that these numbers may well have been surpassed.

Governmental use of servers alone can take up 10 – 12% of a nation's annual datacentre electricity consumption[v]. Figures for the private sector are harder to come by. According to other figures presented, datacentres are estimated to consume some 2% of global resources per year with an annual compound growth rate of about 10 - 15%. A single "internet-scale" datacentre[vi] may well have a consumption of some 50 MW, equivalent to some 40 000 homes, typical large datacentres are reported to consume the equivalent of 15 – 20 000 homes.

While energy consumption and thermal dissipation are the principle sources of environmental impact arising from the construction and running of datacentres, impact from the use of water for cooling; from the carbon footprint both of the centre itself as well as its infrastructure and employees; the impact of construction of the site itself must also be considered. We must consider the whole resource efficiency of the centres if we are to minimise their costs and impact.

The cost of construction of a major datacentre is several billion Euros. This budget is presently divided, 40% to IT Equipment, 50% to cooling. Recent volatility in the unit price of power has meant that total cost of ownership of datacentres has grown dramatically over the last few years, even for large-scale consumers.

The Workshop heard that it is possible, at least in the short term to reduce power consumption by an overall 20-30%. This has led to a series of initiatives of which the present one is a part.

One of the most prominent has been by the United States which has been looking into this issue for several years. The US Government, through the Environmental Protection Agency (EPA)[vii] has been evolving a series of actions based on a mixture of guidelines, best practice, legislation, monitoring and technological advance with the aim of controlling and in the longer term to reduce the rate of consumption of power by datacentres and other large IT consumers. The cumulative effect of these actions, the EPA projected, would be to lead to a reduction in consumption of some 25% with a concomitant reduction in the national energy bill by 2011[viii].

US Public Law 109-431 (2007) states that "It is the sense of Congress that it is in the best interest of the U.S. for purchasers of computer servers to give high priority to energy efficiency as a factor in determining best value and performance for purchases of computer servers." In other words there is also a politico-economic imperative to involve users, producers and all other stake-holders in finding holistic routes to reduction in datacentre impact.

In both short- and long-term Europe too must balance its need for more and more information with longer term reduction in storage costs and the efficiency with which data can be stored. The need for data storage will increase inexorably with the growth not only of developed but also of developing economies; the only way to address this is by addressing existing problems and developing technologies to further drive down power-consumption.

The future of datacentres is a marriage between the IT sector in all its many guises, the energy engineering sector, end-users, the research community, politics and economics. We have an obligation to future generations to develop a route forward; understanding the best that technology can provide; and in the short term making decisions that will enable us to achieve short-term reductions while simultaneously carrying out necessary research to support future reductions.

This must be paralleled by a programme of activities designed to help the business community understand the importance of reduction in the environmental footprint of their datacentres regardless of size.

The European Scenario

We cannot do without data. All sectors of society and of the economy are reliant upon it. The dramatic growth in the use of data and of the sources, not only in the more traditional sources, but particularly in the social sector has driven rapid growth in the number of datacentres being constructed worldwide. This has lead to a very substantial increase in both power consumption and power dissipation.

Where once data storage was a commercial activity this now extends down to the individual as the creation and storage of information generated by the public at large is a major activity around which many organisations are building new business models and which Web 2.0 actively encourages.

Europe is no exception in this. As the largest single economy in the world, with a population of some 500 million, Europe places great reliance upon data, its collection, marshalling and storage. It is also one of the most highly developed groupings of economies in the world and one of the most highly reliant on data. Today 50% of economic value in developed countries is reliant on data. This may be in the financial hubs of London, Frankfurt or Paris; or companies who are managing their production processes through all its stages whether from Lisbon to Warsaw; or individuals wherever they may be. They all rely on fast, global information access.

It has been suggested that in scientific computing we are in the throws of creating a "Fourth Paradigm"[ix] the datacentre; where the masses of data necessary for scientific discovery are stored. The meeting strongly believes that this should be extended to include data that is mined by commercial enterprises and images as well as other forms of data stored by individuals.

The human race already stores vast amounts of data, at present about 1.2 zettabytes (1 zettabyte = 10^{21} bytes), or the equivalent of about 100 newspapers for every person on the planet[x xi]. The rate of growth has been accelerating over the past few years and there is no sign of any diminution[xii]. It is estimated that Western Europe accounts for about a third of this figure.

Governmental and similar organisations store vast amounts of data and this is growing they are also among the major users of computer systems. Financial institutions also regularly store multiple copies of every trade made in every global market. This alone has been and continues to be a major driver of data storage and one crucial to the economic well-being of Europe.

Day-to-day applications generate large amounts of data, one of the largest BEIN image data. The amount of security video held (even if only temporarily) has grown at an astounding rate in many countries, most notably the UK which is probably the world's most surveyed country per capita[xiii]. Security data is supplemented by, possibly surpassed by, consumer video services such as Youtube and other sources. It is often said that modern processing and storage technologies are driven by the need for video storage. In addition there is a vast amount of telephony and related data (calls, tweets, SMS etc.) stored, as are emails and data from a host of other sources.

While the cost of data capacity (as measured in Euros per Megabyte) is dropping[xiv] rapidly the total volume of data is growing at an even greater rate. Looking to future applications, such as those involved in Smart Grids, Systems of Systems and similar complex applications, it is easy to understand why an even greater data deluge is predicted.

Clouds, whether private, public, hybrid or other varieties have also increased reliance on data storage and accessibility and rely on the existence of high performance, high-volume datacentres. Their growth has been and is predicted to be, rapid. According to IDC global revenue from public IT cloud services grew a base of about $16 Bn in 2009 and will reach $55.5 Bn in 2014, a CAGR of 27.4%, five times that of traditional IT (5%). The economic downturn has actually enhanced cloud adoption due to companies' desire for cost-cutting. Whereas 70% of global revenues was in the US in 2009, by 2014 this will be 51% with Western Europe and Asia/Pacific growing rapidly[xv]. These figures exclude private clouds, which are also growing rapidly. Cloud is expected to be a major driver in the growth of datacentres.

Finally we must recognise that we cannot anticipate future uses of data and the means by which it might be generated. One has only to look at the growth in mobile telephony, social networking and imaging systems of all kinds over the last decade to appreciate that.

Europe has a major role to play in reducing global environmental impact resulting from industrial growth, which it is addressing through the 20-20-20 objective of the European Strategy 2020, namely (with reference to 1990 levels) to reduce greenhouse gas emissions by at least 20% (30% if conditions permit); increase the share of renewables in final energy consumption to 20%; and increase energy efficiency by 20%. Datacentres have a significant role to play in achieving this goal.

Europe has a unique capability to build on its world-leading expertise in building low-cost, low-power systems that will play a major role in achieving these goals, on a global as well as a European stage. Companies such as ARM, ST Microelectronics and others already create state-of-the-art low-power systems for resource-constrained systems such as mobile phones. What is now needed is to bring that expertise to bear in creating resource-efficient, low-power systems for the datacentre.

Measuring Impact

Critical to controlling a datacentre's environmental and economic impact is the ability to measure its power consumption and its thermal output.

As a result of the industry's desire to understanding these better, a great deal of effort is being expended in seeking to agree appropriate criteria and metrics. There are efforts, particularly among the CIOs, CTOs and other board members through organisations such as The Green Grid[xvi] and CTO Alliance-CRIP[xvii]. The Green Grid is a global network of businesses and executives who have set out to create a suite of metrics that can be deployed and are readily understood at board level. The CTO Alliance-CRIP is a European non-profit organization gathering CTOs, chief architects and senior IT infrastructure managers from major European corporations and public entities, who have been working since 2007 on datacentres, energy efficiency standards and operational best practices.

Globally activities are also underway by several standards bodies to create metrics including at the ITU[xviii], ETSI[xix] and jointly by ISO and IEC[xx]. In order to achieve global comparability there will need to be a harmonisation of approach among metrics if they are to be used to monitor what is a global issue.

Green Grid have evolved PUE[xxi] and DciE as measures of datacentre efficiency. PUE or Power Usage Effectiveness is a measure evolved to determine the overall energy efficiency of a datacentre; DciE ("Datacentre infrastructure Efficiency") is the reciprocal of PUE.

PUE is determined by dividing the amount of power entering a datacentre by the power used to run the computer infrastructure within it. Alongside PUE The Green Grid alliance also proposes the use of Water and Carbon Usage Effectiveness in order to get a better perspective on such systems.

As metrics PUE and its relatives are attracting attention from senior executives because of the clarity of their approach. Furthermore this interest is global, so figures can supposedly are comparable across the industry.

Nonetheless there is controversy regarding the applicability of these high-level tools for anything more detailed than an overview of consumption. Those taking this view seek to create more detailed metrics and techniques that enable the managers of datacentres to provide clear and scientifically valid measures of environmental and economic impact. While PUE and its relatives provide a high level perspective, it must be realised that among the component parts of very large corporations, and between corporations there can be large variations in technology and in circumstances, which restrict the validity of such an approach.

The Workshop welcomed PUE as a first step and discussed the need for the evolution of more detailed and locally applicable technologies and metrics to better control and monitor datacentres.

To maximise the benefit to be gained from improved metrics we require tools that permit the accurate and effective measurement of consumption, dissipation and cost-efficiency across a spectrum of parameters. This should form the basis of an ongoing programme reflecting not only changing usage patterns and technologies, but also taking into account the differing needs of the various strata of industry.

Datacentres are not “one-sized”. Those employed by the likes of Amazon, Google and similar global, information-based organisations is vastly different from those used by SMEs, who may rent a small part of a larger system or build their own, while those of mega-corporations may cover many hectares and have their own power supplies.

Many very large information-driven companies are often more conscious of the impact that they are having on the environment than smaller ones. Smaller corporations tend to monitor their consumption but, for reasons associated with their structures and accounting practices, often fail to take action to control usage.

However, while an SME may often save only a few percentage points of their total costs by taking action to make their datacentre more efficient, SMEs dominate the European economy[xxii]; there are far more SMES than there are multi-national corporations, and were a substantial proportion of the medium-sized and small SMEs together with a good portion of micro-SMEs to be assisted to reduce their power-consumption, it would have a disproportionately large economic and environmental impact.

For these reasons the Workshop considered that it was important that action be formulated to increase the efficiency of datacentres for all companies, and that as part of this action, assisting SMEs across Europe to improve their datacentre's efficiency should be a major aim.

Incentives To Reduce Cost and Environmental Impact

It is a myth that businesses of all sizes are not interested in being “green”; they can be incentivised to become green in many ways. It is also true to say that many industries are going green of their own volition driven on by the imperatives of running any business, including the profit motive, but also including their awareness of social issues among their communities.

According to a survey quoted in “The Economist” some 83% of companies surveyed prior to the Durban Round of the UN Climate Summit wanted to see it generate a legally binding agreement to put a price on carbon emissions[xxiii]. Moreover firms are making green investments because they see these as making sense. The same article states that environmental strategies are central to a growing number of major companies. These include investments in energy efficiency and/or renewables, the majority of which will rapidly pay for themselves. The largest data-oriented companies are among the most efficient according to information presented to the Workshop.

Whilst economies of scale, benefit to reputation and/or image may well reward very large organisations, this may not be the same for the middle-ranking companies let alone for SMEs. Their business issues are different. Nonetheless the realities of the future are such that they too will have to become greener.

The Workshop considered that incentivisation of businesses should be undertaken through convincing case studies that demonstrate the business value yielded by improvements in efficiency. These should show clearly that such benefits go far beyond cost savings. The studies must be crafted in such a way as to clearly address the concerns of all the various strata of industry.

The workshop considered inter alia, the UK's CRC Energy Efficiency scheme which targets larger consumers[xxiv]. This is a mandatory scheme aimed at improving energy efficiency and cutting emissions in large public and private sector organisations responsible for around 10% of national emissions. The scheme features a range of reputational, behavioural and financial drivers, which aim to encourage organisations to develop energy management strategies that promote a better understanding of energy usage[xxv].

As a result of this and other experiences it was felt that an attempt to regulate directly or legislate for greater efficiency on a European scale would be impractical and take a great deal of time to implement.

The Workshop felt however that there are approaches which offer easier ways in which to address the problem and which can reach most companies that use datacentres. Briefly these include:

- Making the resource-efficiency of datacentres clear to all and a part of the costs to all users (transparency)

- Encouragement to adopt commercial measures that will encourage datacentre service providers to deliver cost-efficient data storage to their customers.

- The adoption of technologies that will enable increases in the resource efficiency of data-intensive industries, whether datacentres or others. These include processor and other hardware technologies (including storage media); Power infrastructure and Power Supply Unit (PSU) technologies; cooling technologies; and data delivery technologies including internet.

- Programmes to promote awareness among industry of the benefits of improved efficiency.

Many of these also feature in US Public Law 109-431 (see above), however from a technological and economic point of view Europe is probably better placed to promote such activities, due its greater coherence and the rapid evolution of the IT sector among the more recently accessionary states among others. Europe's exceptionally strong position in low-power technologies also provides an opportunity for it to become a world-leader in reducing the environmental and economic impact of datacentres.

Today a considerable cross-section of industry outsources its data storage requirements. While the exact percentages are unclear, the provision of data storage as a service means that business pays integrated charges. The Workshop suggested that one effective way of reducing impact would be to promote greater transparency of consumption, by billing not only on an overall figure but also using a range of metrics, such as overall power consumption, PUE and kWh per TB, stored to enable the purchaser to make comparison among suppliers. This would encourage users to shop around to minimise the costs and increase competition among service providers.

Such concepts could also be pursued through terms in Service Level Agreements (“SLAs”). SLAs drawn up with reference to such metrics would provide additional incentives for datacentres to reduce their costs and would broaden market choice.

To accompany SLA measures and other measures to increase transparency there needs also to be investigated measures would enable rapid and easy migration of data should a user wish to move datacentre irrespective of size.

By promoting such concepts, the Workshop concluded, real and effective pressure could be imposed to increase efficiency and reduce costs, without the need to legislate.

The Workshop noted that there are straightforward advances that can be made in the better utilisation of power consumed. Often a new installation will, for economic reasons, not be fully-populated at start-up, nonetheless the centre is cooled as if it were at maximum capacity. Were the centre to be better constructed it should be easy to only have to cool the populated section of the facility. Producing more detailed information on consumption will drive more accurate measurement of heat and power flows and improve efficient heating (and cooling) of centres.

Thus there are obvious “low-hanging fruit” available which can be addressed. Implementing these is expected to lead to a 20-30% reduction in consumption in the short-term[xxvi]. Achieving, maintaining and then increasing such reductions will require investment by all European stakeholders in developing new technologies and in implementing them rapidly, in the short, medium and long term.

Achieving The Goal – Technology

The Workshop had as one of its starting points the output of the June 2011 Workshop on Next-Generation Datacentres (see box below).

With a long history of innovation and major breakthroughs in embedded and mobile platforms, energy-efficient technologies and a corresponding strong industrial backbone, Europe has the opportunity to take the lead in research and innovation to bring 100x improvements in datacentre efficiency and total cost of ownership by 2020 needed to sustain the demand on information technology growth.

New technologies are required for future datacentre hardware. As a reference point, we take the widely quoted figures for a typical server power consumption as given by Fan et al[xxvii] in their review of datacentre architectures. These are widely quoted and have formed the basis of a large number of studies and comparisons. They are based on a single server and don't take account of disk farms etc[xxviii]. Nonetheless they are capable of forming the basis for modelling and comparison.

The figures are as follows:

Component	Peak Power Watts	Percentage of Total
CPU	80	32%
Memory	36	14%
Disks	12	5%
Peripherals	50	20%
Motherboard	25	10%
Fan	10	4%
PSU Losses	38	15%
Total	251	100%

Fig 1. Data Server Component Power Consumption,
Derived from Fan et al *op. cit.*

Power-consumption has changed somewhat and these figures need to be revised, but by adapting them we can understand better relative improvements.

The Workshop identified a number of important areas of technology that should be addressed as priorities if we are to maintain the reduced levels that might be achieved through the “low-hanging fruit” and to assist in minimising and then further offsetting future overall growth. A programmes should have as its target long term improvements by two orders of magnitude and individual projects (to form part of WP2013) should give measurable evidence as to how they will achieve these goals. Such targets are not unrealistic given anticipated developments in server technologies in general and processor technologies in particular.

While the Workshop had outlined ways of achieving short term gains, in order to take the process further requires the creation of new metrics so that accurate measurements can be made. Existing ones (such as PUE) are broadly specified and as a result fail to reflect the detailed structure, loads or usage patterns of actual datacentres. In doing so they do not necessarily encourage best practice at the datacentre management level. PUE therefore needs augmentation by metrics robust enough to show how efficient a real centre is, where inefficiencies arise and to enable managers to build on best-practice.

While these metrics, and their near-relatives for measuring carbon and water consumption, provide broad overall frameworks within which to measure efficiency at a high level, more effort is needed to harmonise the way in which they are measured. At an international level Europe should play an active role in bringing together the various organisations that are already working to this end (e.g. Green Grid, US EPA, US DOE, EU's JRC, Japan's GIPC).

While best practice will stimulate the exchange of information and the evolution of better approaches to monitoring efficiency, the Workshop felt that mechanisms needed to be put in place to enable the creation of local metrics to supplement, and build on, best practice so that individual facilities can address their own local issues in a way that is consistent with industry practice.

Work on new metrics needs to start in the near future; but without a longer term programme of continuing research and development their effectiveness risks being negated in future as data volumes increase exponentially and datacentre technologies change. Future plans must be driven by the need for power-reducing technologies to outperform the growth of data.

Nonetheless a great deal of work will have to be put into exploiting these advances. The following, non-exhaustive, list of areas for focus were suggested (all with regard to the performance goals above):

(i) **Applications Software Efficiency**: datacentres rely on software to manage them and ultimately it is the efficiency of software that controls the number of systems that need to be deployed to achieve given levels of throughput and storage. Efficiency understood here as meaning multi-parameter optimisation, covering such as performance, energy, cost and timing. Modern codes are already complex but the need for data to be stored *and* retrieved *and* analysed with regard to the amount of power consumed will add a further level of complexity. Programming complex systems such as datacentres is already an active area of research, for example in FP7 Project RELEASE[xxix].

(ii) **Hardware/software co-design** is also an area of active research particularly in embedded systems where energy consumption is critical. Low-power systems is an area of semiconductor technology in which the EU leads the way. It has recently begun to be deployed in hardware and systems design for datacentres in an effort to achieve increased efficiency. The Workshop recommended that activities in this area should be undertaken as an important route forward.

(iii) **Advances in processor design** must be exploited in order to minimise the amount of power consumed, for example through the use of additional lower-power cores in multicore systems in order that processes can remain in a stand-by mode that uses very low levels of power[xxx] but nonetheless allows a fast switch-in when required, as well as approaches such as 3-D stacking (see for example: FP7 Project Eurocloud[xxxi]). Europe has a world-lead in the area of low-power systems which should be built upon in order to improve environmental, economic- and resource- efficiency. These technologies can together have a very important impact in “Going Green, Going Lean, and Going Clean”.

(iv) **Data storage and data access** is the *sine qua non* of any datacentre. Therefore investigation needs to be undertaken into new and novel ways of storing data , including disk technologies, targeted at reducing overall power-consumption, increasing performance and throughput and simultaneously reducing storage costs. There is already activity here, such as the FP7 Project IOlanes[xxxii].

(v) **Tools** should be designed that are capable of providing interactive monitoring of datacentre metrics capable of permitting the system's administrators to configure it “on the fly” for increased efficiency, within the requirements of the centre’s parametrisation.

(vi) **Research into federated, collaborative datacentres** which address local accessibility of data and sharing of resources. Related to this is the idea of “chasing cheap energy” where data, or data access, is moved among federated datacentres distributed so as to be able to use low-cost electricity.

(vii) **New approaches to minimising cooling requirements**, to reduction of power consumption and to increase the re-use of heat generated are needed. The Commission already sponsors work on the increase in cooling efficiencies through FP7 projects such as CoolEmAll[xxxiii] and other projects which aim at an integrated approach to increasing datacentre efficiency. The approaches that these projects are creating need to be evaluated in the context of this programme.

(viii) **Improvements to Uninterruptible Power Supplies** (UPSs). These are central to the maintenance of services but they are a heavy cost, economically and, in some technologies, environmentally.

(ix) **Interconnectivity** The Workshop stressed the importance to any programme of reduction in economic and environmental impact of datacentres that they have reliable, high-speed access. Interconnectivity enables the large corporation, the SME and the citizen to access, use and manage their data better. As such connectivity must be seen as a crucial tool to enable global competitivity by a data-driven Europe.

(x) **Renewable energy sources** The Workshop acknowledged that in the short term it appears impracticable to use renewables as a means of powering datacentres, however it was urged that investigation be made into their longer term viability. However there was a clear belief that datacentres should, indeed must, aim to "go Green" in the longer term through improving the ratio of renewables in the power source mix.

The Workshop stressed the need for new technologies to target all strata of datacentres regardless of scale and also the need to ensure that all users became aware of their datacentre consumption as a means to understanding their own environmental footprint.

Workshop on Next-Generation Datacentres – Summary of Outputs

The report from the Workshop on Next-Generation Datacentres (Brussels, June 2011)[xxxiv] analysed the main research challenges in datacentre design from a European perspective bringing together industrial representatives including ARM, Bull, ST Microelectonics, Alcatel, SAP and Ericsson. It defined these challenges to be:

• **Achieving mobile systems resource efficiency in server nodes**: Digital platforms require drastic improvements in cost, resource- and energy- efficiency, to enable them to achieve similar levels to those found in mobile platforms. The report stated that an improvement of at least two orders of magnitude in the energy consumption of server nodes should be the goal.

• **Addressing data volume**: Information technology is changing from being "computation-" to "data-centric". Making sense of massive amounts of data requires a complete redesign of data management technologies and innovation in machine learning and probabilistic computing to reduce the size and dimensionality of data for services and applications.

• **Hardware and Software Integration**: future datacentre designs require collaboration between software and hardware to provide both energy (through specialization) and thermal management especially for future 3D-stacked server chips.

• **Federated datacentres**: Despite the trend towards single or few giant datacentres, security, reliability/availability, affinity to data, proximity to energy sources, legal boundaries and the client base dictate the need for technologies to orchestrate operation and load balancing over a collection of physically distributed sites.

Towards a European Roadmap to Address the Economic and Environmental Impact of Datacentres

We are storing data at an ever-increasing rate and the data storage industry is growing rapidly and consuming increasing amounts of power as it does so. The heat from these activities is then dissipated, ultimately into the environment. There is thus a two-fold economic and environmental impact. In order to overcome this we have to formulate an approach that will firstly match the rate at which data is growing with reductions in cost and impact, and which will go on to reduce the overall costs in the longer term.

This Workshop has made a number of recommendations for urgent action set out in detail above and we recap these below. The combination of impact on the environment; of economic drivers from the perspective of the datacentre users and owners; and the increasing volatility of energy pricing on one hand; and the data deluge means that action must be taken soon.

The goals of the programme which would arise from the recommendations of the Workshop would contribute to meeting both European and global goals on climate change; on reducing environmental footprint; and on energy consumption. Clearly they alsose target the triple goals of a lean, efficient and "green" industry. Because it is global there needs to be integration among efforts and in this the EU through its well-known activities (Codes of Conduct for DataCentres[xxxv] and UPS's, Framework Programme inter alia) and its industrial and technological leads is well-placed to take a leading role.

While component parts can be worked on independently, ultimately any approach implementing these measures will be holistic and cross-cutting, targeting a number of areas in particular: Policy and Economic Drivers; Technology; and Business and Service Models. Only a programme with integrated aims and which draws on many varied resources can address such an enormous problem.

The areas on which the Workshop recommended that action be taken divide into two parts:

(i) **The"low-hanging fruit":** These consist of a number of actions which can be taken to encourage greater efficiency in the use of existing resources.

The meeting noted that different strategies are required to engage different strata of business; those for the very large organisations differ markedly from those for SMEs. Nonetheless the whole spectrum of industry is susceptible to “greening”. The participation of SMEs is important since while individually they may have relatively limited data requirements their number means that their participation in any action would have considerable overall impact.

Such measures should motivate business by demonstrating greater efficiency and business benefit. These can take a number of forms: greater efficiency, greater effectiveness and business edge among others.

These should take place in conjunction with a series of measures which would oblige datacentre service providers to provide metrics to clarify the true costs of datastorage. These might include an obligation to provide PUE and cost per MB stored, among others. In the short- to mid-term further metrics should be developed in order to monitor efficiency and for the better and more accurate monitoring of efficiency.

These outputs should be included in Service Level Agreements and other tools which would increase visibility of the consumption and the impact of individual datacentre service providers. It must also be made easy to move data between centres.

(ii) **Longer term programs** to foster new European technologies with the aim not only of achieving the twin goals of reduction of environmental impact and greater efficiency, but also of positioning Europe as a world-leader in datacentre efficiency and related technologies.

The components of such a programme, which are necessarily complex, in order to create an integrated approach to a complex problem are set out elsewhere in this report and are not further detailed here.

The workshop clearly felt that the scope for primary legislation to encourage datacentre owners and users to regulate their environmental and economic impact was small given the need for rapid impact and the complexities of enacting and enforcing such legislation.

The European Code of Conduct on Datacentres should be enhanced to help address the issues that arise here in two phases: the first to address the short-term opportunities and the second to encourage and enhance the adoption of the technologies developed under the activities proposed by the Workshop.

In creating a plan of action based upon the recommendations of this, and the other workshops and sources quoted, the Workshop recognises that it is asking for the expenditure of a substantial resources. However it is to be borne in mind that the impact of datacentres constitutes a rapidly growing problem which must be addressed. If it is not tackled through a program such as that set out here, the reality is that the impact on Europe and ultimately on the world would vastly exceed any expenditure sought under this roadmap.

Attendees

Rapporteurs

Workshop Rapporteur

Peter Dzwig	Concertant LLP

Session Rapporteurs

Mark Acton	451 Group/ Uptime Institute Network
Andrew Donoghue	451 Group
Andrew Lawrence	451 Group

Participants

Aguirre de Cárcer	Guillermo	Barcelona Supercomputing centre
Bertoncini	Massimo	Engineering Ingegneria Informatica
Christmann	Wolfgang	Christmann informationstechnik + medien GmbH
Davies	Craig	Maxeler Technologies Ltd
De Bosschere	Koen	Univ. Ghent
Di Girolamo	Marco	HP Italy Innovation centre
Dobbelaere	Philippe	Alcatel-Lucent
Falsafi	Babak	EPFL
Gilmore	Mike	e-Ready Building, ETSI STF on Global KPIs for DC
Haas	Peter	High Performance Computing centre Stuttgart
Juan	Ana	Atos Origin
Kennedy	John	Intel
Kenway	Richard	EPCC (Edinburgh Parallel Computing Centre)
Klingert	Sonja	University of Mannheim
Koller	Bastian	High Performance Computing centre Stuttgart
Kolodner	Hillel	IBM Israel
Muggeridge	Malcolm	Xyratex Technology Ltd. —

		Network Storage Solutions
Pernici	Barbara	Politecnico di Milano
Philips	Ian	ARM
Pierson	Jean-Marc	IRIT Universite Paul Sabatier Toulouse
Roche	Dominique	Orange
Rouyer	Andre	Green Grid
Rudgyard	Michael	Concurrent Thinking
Sazeides	Yanos	UCY
Senata	Ahmet	Barcelona Supercomputing centre
Sicsic	Pierre	HP
Stephan	Francois	CRIP (Club des responsables d'infrastructures et de production)
Tong-Viet	Emmanuel	IBM
Treptau	Patrick	1E
Baxevanidis	Kyriakos	European Commission
Chastanet	Pierre	European Commission
D'Elia	Sandro	European Commission
Heap	Tim	European Commission
Riemenschneider	Rolf	European Commission
Tsakali	Maria	European Commission

References

[i] "Technology-Scalable Datacentres" Report from the Workshop on Next-Generation Datacentres, June 2011, Unit G3: Embedded Systems and Control. http://cordis.europa.eu/fp7/ict/computing/documents/nextgen-dc-final.pdf

[ii] European Commission, Code of Conduct on Data Centres Energy Efficiency, Version 1.0, European Commission Directorate-General JRC Institute for Energy, Renewable Energies Unit available at: http://ec.europa.eu/information_society/activities/sustainable_growth/docs/datacentre_code-conduct.pdf

[1] [iii] This is based upon the Draft UK Market Transformation Programme European Enterprise Server Installed Base Model, assuming an upper bound ratio of 1:2 between electricity consumed by the server equipment within the data centre or server room, against that consumed by cooling equipment and through power loss. The lower bound ratio of 1:1 gives total electricity consumption close to 37 TWh. The upper and lower bound ratio is based on several different sources of measurements of electricity consumption in the data centre.

See http://efficient-products.defra.gov.uk/ and http://www.mtprog.com/

According to some models (see ref 4 below and elsewhere) even the upper bound may actually represent an underestimate as data growth may turn out to be higher than expected.

[iv] US Environmental Protection Agency, ENERGY STAR Program, "Report to Congress on Server and Datacentre Energy Efficiency Public Law, 109 – 431, 2007"

[v] See reference 4 above

[vi] This is for a fairly large-scale datacentre providing global access to major data resources.

[vii] http://www.epa.gov.org/

[viii] At present the results are not available, although it is likely that delays in implementation and the impact of financial crises, commodity price volatility and other events over the intervening period will have pushed this laudable goal somewhat off course.

ix Hey, AJG et al “The Fourth Paradigm Data-intens

ive Scientific Discovery”, Microsoft, ISBN 9780982544204, 2009

x As of 2009 the Internet was estimated to contain 0.5 ZB. In 2011 Seagate alone report having sold 0.3 ZB of drive capacity.

xi IDC estimate that the total amount of data in th

e world passed 1.2 ZB in 2010, note that this figure includes non-stored data, although it is dominated by stored data. International Data Corporation, “The Expanding Digital Universe”, 2007

http://www.emc.com/collateral/analys

t-reports/expanding-digital-idc-white-paper.pdf

xii Source IDC, ref 10.

xiii Lewis, Paul. "Every step you take: UK underground centre that is spy capital of the world", The Guardian, March 2, 2009, see: http://www.guardian.co.uk/uk/2009/mar/02/westminster-cctv-system-privacy

xiv The current (Dec 2011) retail price of 1TB is around 80 Euro.

xv See http://www.idc.com/prodserv/idc_cloud.jsp and sources therein.

xvi See http://www.thegreengrid.org/

xvii CTO Alliance France (CRIP) was created in 2007 (http://www.crip-asso.fr/); CTO Alliance UK was created in September 2011; CTO Alliance Luxembourg in October 2011 and CTO Alliance Germany will be created in February 2012.

xviii ITU Study Group 5 http://www.itu.int/ITU-T/studygroups/com05/index.asp

[xix] ETSI Technical Committee on Environment and Engineering http://portal.etsi.org/portal/server.pt/community/EE/303) and Technical Committee on Access, Terminals, Transmission and Multiplexing http://portal.etsi.org/portal/server.pt/community/ATTM/297

[xx] ISO/IEC Joint Technical Committee 1 on Energy Efficient Data centres (EEDC) http://isotc.iso.org/livelink/livelink/open/jtc1/

[xxi] See "WP#22-Usage and Public Reporting Guidelines

for The Green Grid's Infrastructure Metrics PUE/DCiE" at http://www.thegreengrid.org/en/Global/Content/white-papers/Usage%20and%20Public%20Reporting%20Guidelines%20for%20PUE%20DCiE/, October 2009

[xxii] 99% of all European businesses are SMEs. They provide two out of three of the private sector jobs and contribute to more than half of the total value-added created by businesses in the EU. Source: Europaen Commission DG Enterprise and Industry at: http://ec.europa.eu/enterprise/policies/sme/facts-figures-analysis/index_en.htm

[xxiii] "Why firms go green", The Economist, Schumpeter column, November 12th 2011

[xxiv] See http://www.decc.gov.uk/en/content/cms/emissions/crc_efficiency/crc_efficiency.aspx

[xxv] In particular the scheme has been the subject of legal challenges and has had to undergo substantial revision. See previous reference

[xxvi] Based on data in the Energy-Star report to the US Congress and other sources quoted in the meeting.

[xxvii] Fan, Xiaobo Weber W-D and Barroso, L. A. 2007. "Power provisioning for a Warehouse-sized Computer. Proceedings 34th Int Symp on Comp. Arch., San Diego, CA Association for Computing Machinery, ISAC 07

[xxviii] Some indication of the relative importance of disk technologies can be taken from combining the lines "Peripherals" and "Disks" in Fig. 1, while noting that not all peripherals are disk-related.

[xxix] See http://www.release-project.eu/index.php

[xxx] As is already being done in certain devices. See for example: Tegra 3 (Kel-Al series) at: http://en.wikipedia.org/wiki/Tegra

[xxxi] See http://www.eurocloudserver.com/what-is-eurocloud-project-and-our-vision

[xxxii] See http://www.iolanes.eu/about.html

[xxxiii] See http://www.coolemall.eu/

[xxxiv] See reference 1.

[xxxv] See http://ec.europa.eu/information_society/activities/sustainable_growth/docs/datacentre_code-conduct.pdf and also: http://ec.europa.eu/information_society/activities/sustainable_growth/docs/recommendation_d_vista.pdf