European Commission

A Roadmap for Advanced Cloud Technologies under H2020

Recommendations by the Cloud Expert Group

Editors
Lutz Schubert [USTUTT-HLRS]
Keith Jeffery [ERCIM]
Burkhard Neidecker-Lutz [SAP]

December 2012

Digital Agenda for Europe

Reproduction is authorised provided the source is acknowledged.

A ROADMAP FOR ADVANCED CLOUD TECHNOLOGIES UNDER H2020

Recommendations by the Cloud Expert Group

Editors. Lutz Schubert (USTUTT-HLRS), Keith Jeffery (ERCIM), Burkhard Neidecker-Lutz (SAP).

Abstract. The IT industry is facing a growing amount of technological challenges. These barriers affect not only Europe, but the whole world, though with different timescales and different strength, depending on the industrial orientation of the according country. Europe has a great potential to not only face these barriers, but to gain a key position in the global CLOUD market, due to its specific technological expertise and capabilities. The second report on European CLOUD Computing ("Advances in CLOUDs - Research in Future CLOUD Computing", http://cordis.europa.eu/fp7/ict/ssai/docs/future-cc-2may-finalreport-experts.pdf) has identified the set of research areas that Europe can and need to address. This document analyses the expected market development in which Europe can participate and which technological challenges will have to be faced in order to reach these scenarios. It proposes a concrete roadmap to help Europe overcome these challenges and reach a key position in the global CLOUD market.

EXECUTIVE SUMMARY

User and provider demands are increasing faster than supply, and the whole IT industry faces multiple constraints, including performance, green issues etc. CLOUDs provide one way of dealing with this. The European market is highly heterogeneous, dispersed and fragmented, but at the same time it is highly agile and needs to satisfy multiple users at the same time. Together with the high degree of expertise existing in Europe, there is a unique chance for European industry to gain a long-term leading position in the CLOUD domain [SCH12].

Primary business sectors that can contribute to the European leadership in CLOUDs are: Infrastructure provision; Concurrent engineering and simulation; Financial or retail services; Digital entertainment; Government-CLOUD services; Data storage and data centric services; Technology provisioning; Consultation and associated training & education; and various other specific service niches.

The expert group has identified the 10 main technological topics that need to be addressed in order to create the necessary advance in the CLOUD environment which are detailed in the second expert group report [SCH12]. We can classify the R&D topics into immediate, sustainability and game-changing topics:

Immediately relevant work includes: Managing the data deluge; intelligent networking; elastic applications; performance and portability; vulnerabilities; reducing lock-in; competition and collaboration; viable business models;

"Sustainability" actions are: efficient application and user / usage behaviour, including enhanced metadata for autonomy and advanced software engineering; dealing with the growing heterogeneity within and across multiple CLOUDs; scalable orchestration of services across CLOUD platforms to provide an appropriate execution environment including trust, security and privacy;

"Paradigm shifter" tasks include: managing big data (integrity, state); programming style and technique; interoperation and federation; exploitation of new device types and environments;

All of these types must be addressed at least partially right away, as they each have different gestation periods, due to their complexity and technological pre-requisites.

CLOUDs are thereby not only global, but also multi-disciplinary, thus necessitating close collaboration between multiple areas. At the same time, the results will be beneficiary to multiple domains. It also means that the outstanding work can base on a solid foundation of research and development, that should be taken up and exploited. International collaboration can extend and compensate the European situation by offering additional expertise and resources. Europe should thereby proceed in a fashion of collaboration wherever beneficial, but be wary to maintain competitiveness.

TABLE OF CONTENTS

INTRODUCTION & BACKGROUND

Two major trends are changing increasingly fast: The first is the expectation and demand for ICT by end-users, both private and in business. The second is the availability of technology.

The Gartner hypecycle curve indicates that CLOUD Computing is entering the trough of despond before gaining widespread acceptance in some years' time [FEN11]. This means that there is a current tendency to invest less effort and interest into CLOUDs, which gives Europe a unique chance to leap ahead, defining the future of CLOUDs and the related market players.

The market for public CLOUD infrastructure and platforms is currently dominated by the USA although European Telecommunications companies are entering this market. The proprietary nature of these offerings poses serious barriers towards an open market. However, it is possible to imagine an open market in software services which are interfaced to the various CLOUD platforms – both public and private. Other regions of the world are dealing with the two trends mentioned above by incremental steps following the evolution of current practice. This approach will not be sustainable given the incremental increase in demand versus the slow increase in technology capabilities. It also will likely lead to solidify the current unbalanced position of market participants.

Europe is characterized by a heterogeneity of culture and business practices. It also has an agile SME sector, with companies that often are world leaders in their specialties and are willing to take risks. This must be considered an opportunity, rather than a disadvantage as it forces the European industry to think beyond homogeneous infrastructures with a sufficient amount of resources. Therefore, Europe faces an historic opportunity to 'leapfrog' other world regions and develop the technology, legal and regulatory framework and associated business to play a key role, in the international CLOUD computing market.

This document attempts to breakdown this overarching goal into intermediary steps that not only help to reach this long-term vision, but more importantly help to boosting European's IT & CLOUD industry on the way towards this vision.

A. ABOUT THIS DOCUMENT

All content of this document was gathered through a series of working group meetings between invited experts from industry and academia. Main focus of these meetings was to relate the research and development topics identified in the second CLOUD Expert report [SCH12] to the expected developments of and around the European market. For more details on the specific topics and the position of Europe in the CLOUD market, please refer to the first [SCH10] and second report [SCH12].

The document is structured in 3 main parts: Chapter I elaborates the European business environment and Europe's potential to contribute to future CLOUD business cases. Chapter II identifies the core technological research and development tasks that Europe (can and) needs to address in order to realise these scenarios. Finally Chapter III concludes with a quick summary of the main statements in this report.

1. ACKNOWLEDGEMENTS

This report was made possible by the European Commission; particular thanks go to Maria Tsakali and Ken Ducatel for initiating and supporting the effort. The report could not have been realized without the excellent contributions from all workshop participants.

2. LIST OF CONTRIBUTING EXPERTS

Rosa Badia (BSC); Francis Behr (Liberty Surf); Tim Bell (CERN); Ivona Brandic (TU Wien); Thomas Bonhert (Zurich University of Applied Sciences); Manuel Carro (IMDEA); James Clarke (TSSG); Tom Cool (Sony); Thierry Coupaye (France Telecom Orange Labs); Ernesto Damiani (Università degli Studi di Milano); Frederic Desprez (Inria); Beniamino di Martino (Seconda Università degli studi di Napoli); Peter Dickman (Google); Ake Edlund (KTH); Andy Edmonds (Zurich University of Applied Sciences); Ioannis Fikouras (Sony Ericsson); Mike Fisher (BT); Fabrizio Gagliardi (Microsoft Research Connections); Alex Galis (UCL); Georgina Maria Gallizo Rueda (1&1); Alfred Geiger (T-Systems); Diogo Gomes (IT Aveiro); Andres Gomez (Centre of Supercomputing of Galicia); Yi-Ke Guo (Imperial College London); Thomas Haeberlen (ENISA); Eric Henault (EuroCIO); Giles Hogben (CLOUD Security Alliance); Geir Horn (SINTEF); Keith Jeffery (ERCIM); Robert Jenkins (CLOUDSigma); Ricardo Jimenez (UPM); Bob Jones (CERN); Bastian Koller (Universität Stuttgart); Nectarios Koziris (CSLab NTUA); Chatterjee Kumardev (Thalesgroup); Peter Kunszt (SyBIT at Systemsx.ch); Domenico Laforenza (CNR Pisa); Frank Leymann (IAAS - Universität Stuttgart); Bernino Lind (CLOUDSigma); Ignacio Llorente (Openbula); Guy Lonsdale (SCAPOS AG); Sergio Maffioletti (SyBIT at Systemsx.ch); Andreas Manieri (VENUS-C Project Director; ERINA+ Project Manager); Philippe Massonet (CETIC); Francisco Mingorance (cabinetdn); Silvana Muscella (Trust-IT Services Ltd); Burkhard Neidecker-Lutz (SAP); Timo Nyberg (Aalto University); Karsten Oberle (Alcatel-Lucent); George Pallis (University of Cyprus); Will Parker (Dell); Mark Parsons (EPCC); Dana Petcu (University of Timisoara); Yrjo Raivio (Aalto University); Jonathan Sage (IBM); Harald Schoening (Software AG); Lutz Schubert (USTUTT-HLRS); Ulrich Seldeslachts (LSEC); Georgios Smaragdakis (Deutsche Telekom Laboratories; Technical University of Berlin); David Snelling (Fujitsu-Siemens); George Spanoudakis (School of Informatics; City University London); Charlotte Thornby (Oracle); Abdellatif Touimi (Huawei); Theodora Varvarigou (SCES NTUA); Nick Wainwright (HP); Ian Walden (Queen Mary University of London); Walter Waterfeld (Software AG); Philipp Wieder (GWDG); Yaron Wolfsthal (IBM); Inigo Wilson (Orange Groupe France Telecom).

I. VISION

"Prediction is very difficult, especially about the future".
Niels Bohr

Users increasingly expect to extend their home and personal environment beyond a single machine, to be able to carry their applications, data and configuration settings with them wherever they go and access and use them anytime with any device. Not only does this imply that a user can "carry his own environment", but it also means that the user wants to be able to bring and use his own device at work (BYOD), i.e. to use the working environment without disruption on his own devices and mode of working.

The providers and developers behind such a vision must be enabled to easily create fully portable applications that make use of the whole infrastructure environment according to the best capabilities, without confinement to the user's device(s). The capability to compose services and resources must come naturally at this level, allowing the easy sharing and exchange of data, applications etc. whilst maintaining full accountability, as well as addressing of security, privacy and regulatory concerns.

This vision implies multiple technical (as well as regulatory and policy) challenges, not all of which can be addressed within a short time frame. Instead, this development must be broken down into a series of milestones that are not only technically feasible, but at the same time offer advantages in the short and medium term for the European industry that are – directly or indirectly – marketable.

This section attempts such a break-down:

A. KEY CHANGES IN THE ICT ENVIRONMENT

The IT environment (meaning market, research, industry etc.) is subject to constant changes. However, in particular over recent years, this environment has undergone severe paradigm changes the effects of which are only slowly but surely felt and comprehended (see e.g. [WSH12]). Such substantial changes originate typically from the conflict between technical constraints and user demands, i.e. through the attempt of compensating the deficiencies to the best capability.

It is therefore of utmost relevance to try and identify the key changes to be expected in the next 5-10 years that may or will affect the full IT environment:

- ~2014: new resource / processor types arise, potentially first prototypes beyond Turing[1]
- ~2015: everyone goes mobile, BYOD (bring your own device)

[1] Turing and – to a degree – von Neumann imply a mostly sequential execution with immediate data access (cache, registers etc.). This model is no longer appropriate for the degree of parallelism and distribution reached and the technological challenges faced. FPGAs show a tendency towards going "beyond Turing"

- ~2015-2018: more and more countries join the social network and the CLOUD
- ~2016: personalised devices & services, carry your own environment
- ~2016 intelligent multimodal user interfaces commonplace
- ~2016 intelligent virtual networking over physical networking
- ~2016-2020 almost all software offered packaged as services
- ~2018: most applications use big data; the typical size of data shared between devices exceeds the bandwidth of the network topology by far
- ~2020: first exascale machines (are at least aimed for)
- >2020: new computational models take over

B. EXPECTED BUSINESS DEVELOPMENT AND OPPORTUNITIES

The public CLOUD arena at infrastructure (and to some extent platform) level is dominated by large American players who can offer CLOUD services based on large, already amortized capacity to offset the cost of otherwise under- or un-used resources. Thus, their substantial capital expenditure is already balanced by paying customers. As opposed to this, the European computing / data centre infrastructures are fragmented and split up between multiple small to medium enterprises (with the exception of the telecommunication industry) which exhibits a reversed picture compared with the USA. Rather than brute-forcing an European inroad by just copying the US approach verbatim, we propose a strategy that leverages the strength of the European telecommunication sector to enter into the CLOUD market as an infrastructure / platform provider, and investments into dealing with the heterogeneity and segmentation of the European infrastructure.

Whilst the need and usage of CLOUDs will continually grow, Europe has to find a unique market positions, based on its strengths and capabilities, leveraging the opportunities arising from this development. CLOUDs are important to Europe for three major reasons (cf. [KRO12]):

1. they provide a means for industry, especially SMEs, to utilise more cost-effective IT thus gaining commercial benefits.
2. they provide a means for industry, especially SMEs, to access more advanced ICT than through usual architectures allowing more agile use of computing, and to move faster into new markets with less financial risk.
3. they provide an opportunity for IT providing industry, especially SMEs, to offer their services in an open, global marketplace with few regional constraints.

We can therefore identify the following primary business sectors contributing to the European leadership in this domain:

1. Infrastructure provision for the CLOUD that is highly mobile, flexible and adaptive, through the telecommunication industry, as well as federated groups of cooperating SMEs.

2. Concurrent engineering and simulation pilots in Europe's industrial lead sectors such as automotive, energy, aerospace.
3. Financial or retail services where Europe's high security standards and diverse landscape can provide a new challenge for Europe's systems integrators.
4. Digital entertainment to build on Europe's multimedia sector strengths, test new metadata systems, rendering and integrating different media with highly immersive user experience (3D augmented reality media), developing flexible licensing concepts and promoting multiple device use across platforms.
5. Government-CLOUD services where the modernisation of the public sector can act as a catalyst for open "big data" and open CLOUD solutions. The European Innovation Partnership on Smart Cities should be a laboratory for such experiments.
6. Data storage and data centric services hosted by and for Europe, attracting companies, governments and citizens alike, where users can fulfil their security needs for their data and processing in a usable and compliant way;
7. Technology providers, enabling the realisation and provisioning of future, sustainable CLOUD services.
8. Consultation and associated training & education, not only to create the next generation of leading CLOUD enthusiasts, but also to support new entries in moving their application to the CLOUD, supporting exploitation of it etc.
9. Various service "niches" in which Europe already has a leading position, such as health services, travel agencies, logistics, production management, energy supply, agriculture etc.

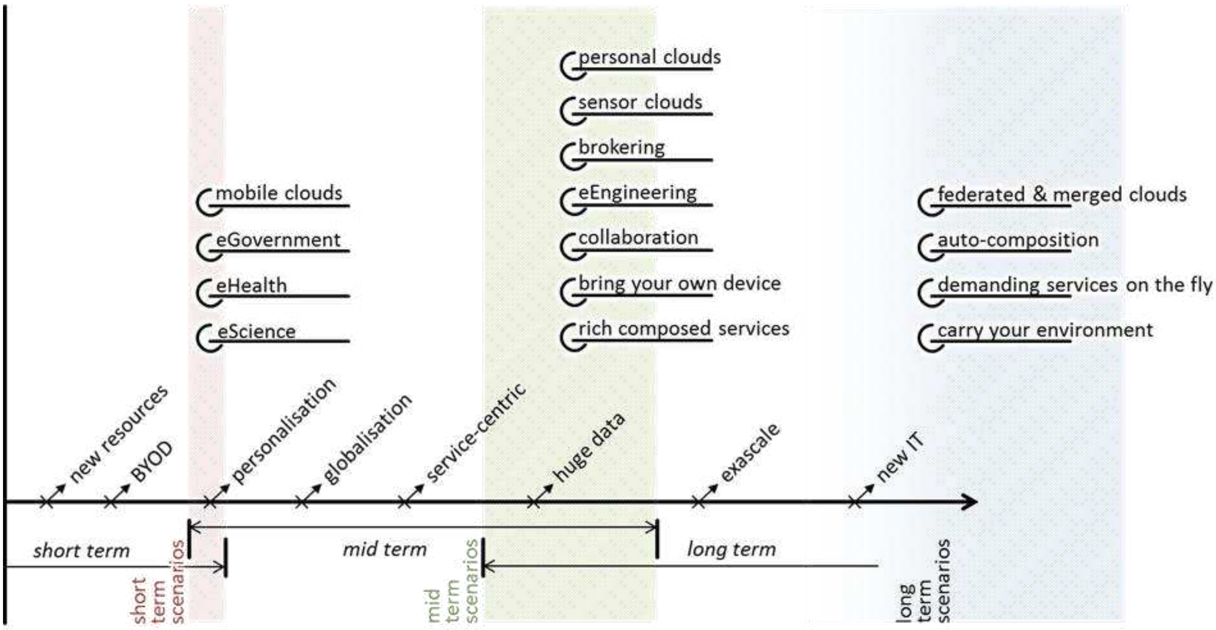

FIGURE 1: OVERVIEW OVER THE BUSINESS SCENARIOS (CIRCLES) AND THE ENVIRONMENTAL CONDITIONS (CROSSES ON X-AXIS). ENVIRONMENTAL CIRCUMSTANCES THEREBY DENOTE "FIRST OCCURRENCES", NOT THAT THEY ARE FULLY FULFILLED

Given the expected development in the IT market and the growing expectations from the users, as well as providers, we can foresee in particular the following use cases and business scenarios to arise over the next years (cf. Figure 1):

1. SHORT-TERM INNOVATION

The European CLOUD industry offers a serious potential for extension and growth in the global market, given the availability of some more advanced technologies (cf. analysis below). In particular we foresee the following potential business scenarios that the European industry may have a good starting position in:

Mobile CLOUDs: Mobility, and the use of mobile devices for personal and office usage will continuously grow over the next 10 years[2]. There is a huge opportunity in services and app provisioning and usage over mobile devices which is not constrained to the US market, but fairly distributed over US, Europe and Japan[3]. Europe must be able to contribute to this mobile growth to remain competitive.

eScience: Scientific simulations such as protein folding, some particle physics, some climate models etc. do not only have academic relevance, but also serve direct industrial needs, such as for medicine development[4]. Embarrassingly parallel problems[5] are ideally suited for CLOUD infrastructures, yet the support is limited as of now. Notably, no supercomputing performance for tightly coupled applications should and can be expected by the CLOUD infrastructure as of now – this does not mean, that supercomputers cannot be exposed *through* the CLOUD.

Moreover the availability of data from any field (from home and city sensors, shared databases) helps to spur interdisciplinary results to emerge (e.g., Nutrogenomics[6]) posing new requirements in term of computational power and data access. CLOUDs can help to handle the ever growing data volumes to be analysed. New distributed data infrastructures will give European innovation the necessary edge – not only in eScience.

CLOUD based services to the ERA communities can thereby help to build and grow relevant expertise. Already on-going initiatives such as Helix Nebula provide a strong starting point in Europe for such development.

Governmental & Social Services over the CLOUD: Maintenance and provisioning of public data and services in governmental sectors, such as education, health care etc. grow in relevance as the international boundaries open. This also involves health warnings, emergency warnings etc. CLOUD infrastructures offer an affordable approach to generate the necessary outreach and availability. Due to their nature,

[2] http://softwarestrategiesblog.com/2012/01/17/roundup-of-CLOUD-computing-forecasts-and-market-estimates-2012/
[3] See e.g. http://mobithinking.com/mobile-marketing-tools/latest-mobile-stats/d#mobilebehavior
[4] See e.g. http://www.biochem.mpg.de/pressroom/Archiv/2011/052_hartl_wieland_preis.html
[5] Problems that can be split up into a number of parallel tasks with little to no problem and implicitly little performance loss through communication overhead
[6] http://www.eufic.org/page/en/nutrition/nutrigenomics/

governmental services are defined by local policies – Europe must therefore develop its own approaches to reach citizens, immigrants and emigrants all over the world.

eHealth CLOUDs: Technical solutions for health as Computer Aided Diagnosis, Personal medicines (as those based on genomic analysis), or treatment or surgical simulations, and Personal Health Services (as storage of patient images) are suitable to be provided on CLOUDs, reducing the cost barriers to access these services. Europe is leading many of these technologies, including initiatives such as Virtual Physiological Human.

Media CLOUDs: Massive growth of multimedia and video traffic will create a need for real-time, distributed media processing in the CLOUD. The information retrieved from data will be consumed and processed in many places concurrently. Handling the predicted growth of video and multimedia traffic is one of the primary challenges that must be addressed in future generations of networks and CLOUDs. European telecommunication industry offers the necessary communication technologies as a substantial starting point.

a) Pre-Requisites

The following technological capabilities, respectively environmental conditions have to be met in order to realise the business scenarios listed above:

- Mobile support including basic support for BYOD
- Reduction of vendor lock-in
- Improved means to deal with data and communication
- Better understanding of CLOUD capabilities and programmability of CLOUD services
- Useable management of trust, privacy and confidentiality
- Adequate market regulatory frameworks and viable business models
- an affordable, international mobile data service

2. Mid-Term Innovation

Mid-term development on the CLOUD market is crucial to ensure the long-term sustainability of the CLOUD industry in Europe. This development is not only defined by the technological changes, but in particular by the arising changes and challenges in the global market in terms of user behaviour and user expectations. At the same time, these changes and expectations give rise to new opportunities, such as:

Personal CLOUDs: Mobility and personalisation of services and environments will constantly increase in relevance over the next years. Maintenance of the personal space and integrating personal devices into a CLOUD like environment arises specifically from the needs of trust and privacy, but also from exploiting existing resources. This is also a particular approach to reduce the communication overhead over the internet (cf. home networks). Europe has a strong position in embedded computing and the telecommunication industry. The interaction of personal CLOUDs with each other, and with (or via) Public CLOUDs, will require significant developments.

Sensor CLOUDs: (related to personal CLOUDs). The future of services is not constrained to "smart" devices, but requires integration of sensors, too. This ranges from smart, controlled home environments, over smart energy grids, up to global emergency predictions etc. Again, Europe has a strong position in both the embedded computing and communication industry. Also, many key smart grid players reside in Europe, as do scientists for emergency situation prediction, such as earthquakes.

Brokering: Reselling and mediating CLOUD infrastructures and services is a growing market[7], as it helps users in making educated choices for selecting the right CLOUD for specific tasks. It also enables SMEs to profit from the large market of providers without undue effort to identify and acquire appropriate resources– this is a particular opportunity for the European heterogeneous market of SMEs.

eEngineering: The distributed nature of the complex engineering projects in the selected sectors and the need of competitive tools (simulation, data analysis, complex visualizations) and the support to SMEs and small Engineering firms are making this market a source of fast innovation in CLOUD technologies and services. Large companies are providing remote engineering tools to their offices around the world, storing their project information in a single site. Spreading the model to small engineering and manufacturing companies providing them with CAD, CAE and PLM tools remotely is a great opportunity for Europe. Small European CAE software vendors can also improve their market share.

Collaborative environments: Collaboration over the internet is the natural evolution of the social network and the growing tendency towards mobile working environments. The shared infrastructure offered via CLOUDs provides the right platform to address this development. The European (service) industry was always centred around B2B and has an accordingly rich set of expertise and experience in collaboration techniques and technologies.

BYOD: The idea of a user's portable mobile device acting as a client-side proxy for the user in terms of authorities and responsibilities is not new but becomes increasingly important in a CLOUD environment. The user preferences and bona fides are stored on the personal device and utilised whenever connected to the internet, including CLOUD-based services – whether a private (organisation) CLOUD or a public CLOUD. Building on the strength of the security infrastructure of the European telecommunication industry, an adequate evolution from the simple

[7] http://www.forbes.com/sites/kevinjackson/2012/08/12/CLOUD-management-broker-the-next-wave-in-CLOUD-computing/
http://www.gartner.com/DisplayDocument?doc_cd=226656&ref=g_noreg

carrier-subscriber relationship to a CLOUD-adequate model can contribute strongly to realisation and exploitation of this scenario.

Enriched services through composition: In an ideal environment services are atomic so that they can be (re-)combined dynamically as required; the (re-)composition being controlled autonomically,with middleware matching the user requirements to available services, data, and other resources. Service offerings need to improve over time to maintain competitiveness. One important approach consists in the combination of individual services and offerings to create enriched capabilities, such as traffic information on the navigation system. Similar to brokerage, SMEs can greatly benefit from the capability to combine offerings for re-selling enhanced services. Europe has already shown its potential in selling enriched services, in particular in the automotive and mobile industry.

One specific example where enriched services should provide great benefits from European industry is Augmented Reality. Europe has technology strengths in media capture, transmission and rendering technologies and research on CLOUD based rich media integration would strengthen the CLOUD computing industry for media applications.

a) Pre-Requisites

The following technological capabilities, respectively environmental conditions have to be met in order to realise the business scenarios listed above:

- Reduced costs
- Increased security
- Performance & reliability, programmability
- Overcoming heterogeneity
- Standards / Interoperability / Orchestration

3. Long-Term Innovation

In the long-term view, CLOUDs will have developed into a commodity, that is taken for granted with most service and infrastructure usage. The market will continue to build up on these capabilities to offer services and capabilities that go beyond the current constraints of CLOUDS. This includes e.g.:

Federated / merged CLOUDs: the IT and in particular the CLOUD environment will continuously grow in heterogeneity and distribution, as will the user and usage scope. To continue offering quality compliant services, mechanisms need to be provided that enable incorporation of resources and services independent of their location or infrastructure, as the individual provider will not be able to match all constraints anymore. Federation and merging involves matching and mapping of the metadata descriptions to determine how the federation may be achieved and at what level of sophistication. A key aspect is the determining of the metadata architecture and making it a standard so that there is a widely acceptable interfacing level for service providers. With the heterogeneity of European infrastructures, Europe has the potential to address this issue well ahead of time.

With federated CLOUD infrastructures, data provisioning and availability, for various big data sectors, including eHealth, eScience etc. can be improved, too, by providing multiple sources. This can at the same time reduce cost for data transfer, in particular for smaller providers. Europe has made considerable advances with this respect in data intensive areas, from which it can build up.

Easy / automated composition: as the use cases will exceed the scope of individual CLOUD providers, so will the services increase in complexity beyond simple straight-forward handling. Multi-tenancy, sharing, distribution, location and in particular personalisation all necessitate more automated ways to create and adjust services on the fly. Metadata describing services – with their functional and non-functional characteristics – permits dynamic (re-)composition. However, the properties of the composition may well differ from those of the individual services and furthermore from the composition itself emergent properties may arise. This scenario has a strong potential for SMEs, similar to reselling, brokering, and enriched services. It does require however increased capabilities to understand, define and execute services.

Demanding services created / composed on the fly: the demands towards richer and more capable services, up to supercomputing capabilities will increasingly grow from the user community, including in particular engineers, eScience etc. For particularly demanding applications honouring the SLA / QoS may require exceptional scheduling management. This is likely to include replicate services for extreme parallelism or adding further CLOUDs to the federation to gain more resource. To realise this type of services, completely new approaches to creating and providing services are needed, compensating the technological deficiencies and maintaining performance. Europe's strong drive in eScience will act as an additional driver to realising these needs.

Carry your environment: The end-user having as a 'bubble' her environment on her client-side portable mobile device raises endless possibilities for commercial exploitation of services. The end-user can discover and compose services to meet their exact requirements including non-functional considerations. There are particular issues in maintaining synchronicity, in identification, authentication and managing access permissions. However for the end-user having in a portable and non-proprietary format information (metadata) on services, datasets, resources and other users allows him / her to have a nomadic existence connecting when required with all credentials stored on the local device.

a) Pre-Requisites

The following technological capabilities, respectively environmental conditions have to be met in order to realise the business scenarios listed above:

- Flexibility, heterogeneity, distribution (new techniques for interoperation)
- Performance & reliability, adaptability (new programming and executing paradigms, new devices)
- Common agreement on the legal framework, including data protection etc.
- Common contractual terminology
- Full connectivity (new data management mechanisms)

II. ANALYSIS

User and provider demands are increasing way faster than that of current technological development. Currently, the whole IT industry is facing multiple physical constraints, such as limitations of latency, clock speed and memory sizes, but also in energy consumption, the overcoming of which is a primary goal of most manufacturers and providers. CLOUDs are one way (and almost certainly the best way) of dealing with it.

Multirack, multiboard (especially with graphics processors), multicore hardware architectures (as typically used in CLOUD infrastructure) require – for efficient use – applications software to be configured appropriately. In the past this has been the domain of the applications programmer. However, there is a paradox. For an individual application high performance (and high throughput) requires a specialised vertical stack from user requirement through code functionality, data utilisation and communications optimisation to reach the (at least partially specialised) hardware. For efficient resource utilisation for a mix of applications the infrastructure platform (everything below the application code) needs to be virtualised in some form (i.e. abstracted and automated). The problem is how to inform the platform of the needs of the application and the application of the capabilities of the platform and, based on these sets of information, how to optimise for the service level parameters agreed. The high degree of complexity of such multinode (multiple CLOUD datacentre) architectures poses considerable challenges in this context.

The quintessential capabilities that future infrastructures and applications need to support are thus generally closely related to the key CLOUD criteria, namely: adaptability, dynamicity, elasticity. The key point being, that code, data and infrastructure need to become much more flexible in aligning themselves to each other, so as to compensate optimally for mismatches between the given capabilities and requirements.

A. TOPIC BREAKDOWN

The European CLOUD Expert Group has identified the following main topics that need to be addressed in order to overcome the deficiencies and offer the capabilities needed to realise the business scenarios elaborated above. For more details, please refer to the second CLOUD report of the expert group [SCH12].

1. DATA MANAGEMENT

Handling the massive amount of data that future infrastructures and use cases will bring necessitates fundamentally new approaches in data handling in the first instances. This will however require a long time to achieve, so that intermediary steps and R&D work is necessary to keep up with the economic pressures.

Intelligent and automated data structuring, as well as definition of the relationship between code and data will help better mapping between the application behaviour, the usage impact and the execution infrastructure. This must respect different timing and communication requirements, including real-time across a large set of sensors and consumers, by using more flexible models of code and data deployment / movement. To this end, the types of storage and the types of data need to be understood better (cf. resource description and usage below).

In all cases, solutions must be able to deal with the massive degree of scale in terms of devices and tenants, as well as with the implicit degree of elasticity and impact of mobility. They must also be able to cope with the increasing redundancy and replication of data, implying out-of-date data and inconsistency issues.

2. COMMUNICATION & NETWORK

The degree of scale in terms of users, applications and data does not match the physical capabilities of communication both with respect to bandwidth as well as latency. Along the same line, users will expect increments in terms of access and accessibility non-regarding those deficiencies. This means, that specific R&D work must be invested into overcoming the physical constraints of hardware on the one hand and into compensating the deficiencies from a software side. To this end, mechanisms are required that enable the control of all communication characteristics from software side, which includes isolated monitoring of performance criteria under multi-tenant usage. The work must strongly correlate with code and data management to account for intermediary offline times due to failure and otherwise, and the according impact on consistency, availability etc.

3. RESOURCE DESCRIPTION & USAGE

Resource types diverge increasingly and with that, the capability to exploit them according to application and usage types. To exploit future resources appropriately on all levels, complete new ways of describing the resource in itself, as well as the relationship between model and function, are required. In this context, formal metadata models will be required to encode all the according information.

The resource information must be exploitable on all levels, ranging from the software stack for mapping and composing capabilities, over compilation down to virtualisation and low-level optimisation of execution behaviour. Notably, in particular on the software and the mapping level, there has already been work done in this context by the Grid community – the according results should be taken up, rather than be repeated.

4. RESOURCE MANAGEMENT

Not only the application needs to be adjusted to the infrastructure, but also the resources within the infrastructure need to be adapted to current usage and its requirements. This concerns all aspects related to scheduling, reservation, placement etc.

Notably, many of these issues have been repeatedly addressed multiple times since the emergence of grids, and in particular Virtual Organisations, where services are placed according to availability of resources etc. Scheduling and reservation technologies etc. are all not new to resource provisioning and should therefore not be treated as something completely new. Instead, the most relevant aspects arise with respect to the impact of the high degree of elasticity, dynamicity and multi-tenancy that has previously only been marginally addressed. Rather than new technologies, this area in particular requires new and updated metrics, considering usage, quality of service, compliance to contractual terms etc.

Resource management belongs thereby to the topics that can never be fully solved, but always need to be addressed, in order to keep up with the growing requirements and expectations, i.e. it will belong to the class of optimisation problems (cf. security below), once the fundamental problems have been overcome.

5. MULTIPLE TENANTS

Under "normal" condition, applications "only" have to scale out across large infrastructures with heterogeneous resources, in a way that tries to reduce communication etc. As long as each user gets access to a completely isolated code and data instance, multiple user requests have no specific impact on the application execution, but the same application must be able to run (with roughly the same QoS) on different infrastructures (heterogeneity).

Scaling therefore impacts differently on parts of code and data, depending not only on the application definition, but in particular on user behaviour, sharing requirements, infrastructure constraints, quality parameters etc. The developer must additionally be aware of the impact of multiple user behaviour, sharing etc.

Note that multitenant usage also has an on impact business models, security, QoS maintenance (isolation of monitoring data, isolation of control etc.)

6. FEDERATION, INTEROPERABILITY, PORTABILITY

This goes beyond the heterogeneity problems of the resource description and portability of the application onto different destination platforms: it also needs to address interoperability across infrastructures, (administrative and legislative) boundaries, data structures etc., including understanding of the application behaviour in such cases and so on. This can be addressed in various ways:

- Standardisation
- enhanced portable and interoperable programming models
- Conversion tools: to convert programs (or source code) and data from one domain into another
- Interoperable middlewares and execution frameworks / runtime environments

The selected means should thereby aim at realising a larger (meta) scale infrastructure that allows European SMEs to actively participate in the global CLOUD game.

7. Programmability & Usability

The classical sequential, Turing like approach is insufficient for the requirements and constraints of future infrastructures, user behaviour etc. The objectives are thereby manifold:

- communication: data-flow orientation, reduction of communication, data structuring, pre-fetching etc.
- segmentation: modularisation, distribution, code structuring, dependencies etc.
- multi-tenancy & sharing: user behaviour encoding, dependencies, shared / nonshared identification etc.
- requirement specification: destination characteristics, quality constraints (including timing, sharing etc.) etc.
- "standardisation": common language models for different infrastructures, applications, use cases / domains.
- other

Because of the dynamic complexity of the inherent e-infrastructure, the application needs to be described e.g. by specifying the functional properties explicitly, allowing subsequent optimisation, interpretation and dynamic scheduling to achieve the expected results. Notably, classical Van Neumann architectures may be insufficient for these cases.

Programming models require some form of compiler (high performance, static) or interpreter (less performance, agility, flexibility) to become executable. It is thereby unclear, whether compilers need to be completely restructured, can be extended or require additional conversion tools. It is also unclear, which specific execution variables impact on the performance, i.e. whether, under dynamic use case conditions given by the CLOUD, interpreters may not indeed provide better performance (as recompilation may be required or conversion may be slow).

Initial work is already under way. Every development must face long uptake times and a strong opposition by existing programming communities, proportional to the degree in which the approach differs from currently accepted ones.

8. Security

With the specific CLOUD / future internet conditions, the security concerns might alter slightly, involving aspects arising from multitenancy. This means in particular that more focus rests on authentication, single-sign-on etc.

An additional aspect of growing relevance consists in the potential impact of outsourcing: even if the data is encrypted and the code signed, the data processor will still have to decrypt data in order to enable processing, and thus can principally inspect and manipulate the data. In general, the actual vulnerability of CLOUD systems are neither well understood, and even less communicated to the end-user, leading to a high distrust in CLOUD offering in the first instance. Besides for better understanding of these vulnerabilities, this also calls for new mechanisms to reduce the power of the data processors, including efficient homomorphic algorithms etc.

Security is an "evergreen" on the IT market, as well as subject to constant development, improvement etc. It must also be considered a crucial factor for in particular industrial uptake. Nonetheless, it must be noted that certain security aspect only become an increasing concern as the underlying technologies provide the respective use cases, e.g. the degree of multi-tenancy, hybrid CLOUD support and so on.

9. Business & Cost Models

With all the changes in the IT environment comes a major gap of knowledge and expertise: providers and developers do no longer know how to develop and provide services that (a) meet the requirements and quality constraints, and (b) earn money and are sustainable.

In this context it is not only relevant to build up an according set of expertise for best practices and cost / business models, but it is also relevant to provide all technical means that can support effective execution of these models – this includes in particular aspects along the line of accounting, billing, monitoring etc. Notably, again, a considerable amount of work has already been performed on the technical side in the context of utility computing, grid, web services etc.

To increase the competence (and hence the competitiveness) of entrepreneurs (including entrepreneurs within established companies) there is a need to support this group in the realization of the value of combining CLOUD technology with business models, and agile & lean operations.

10. Political & Legislatory

Success and uptake of CLOUDs does not only depend on technical progress, but also on political and legislative adaptations between the market players involved. Whilst these aspects are considered out of scope of this report, it should be noticed, that to their support and enactment additional technologies in the infrastructure, such as for privacy support, or for location awareness of the data can and should be developed.

B. Dependency Analysis

As has been elaborated in the second report ([SCH12] p.16ff), CLOUDs are basing on a considerable set of technological advances over the last few years that have enabled many of their capabilities. This means on the one hand that a considerable foundation for further advances already exists, as well as on the other hand, that technological development in this domain is by no way restricted to the CLOUD, but affects and is affected by multiple other domains (cf. Figure 2).

Implicitly, and closely related to this, all these domains face similar technical challenges to overcome which is the primary objective in order to guarantee long-term user satisfaction and thus, more importantly, to boost the European industry sector. We hence need to distinguish between the short-term research and development to boost Europe's market position in the CLOUD the medium-term focus to ensure maintenance and the long-term to gain leadership in this market sector (see below).

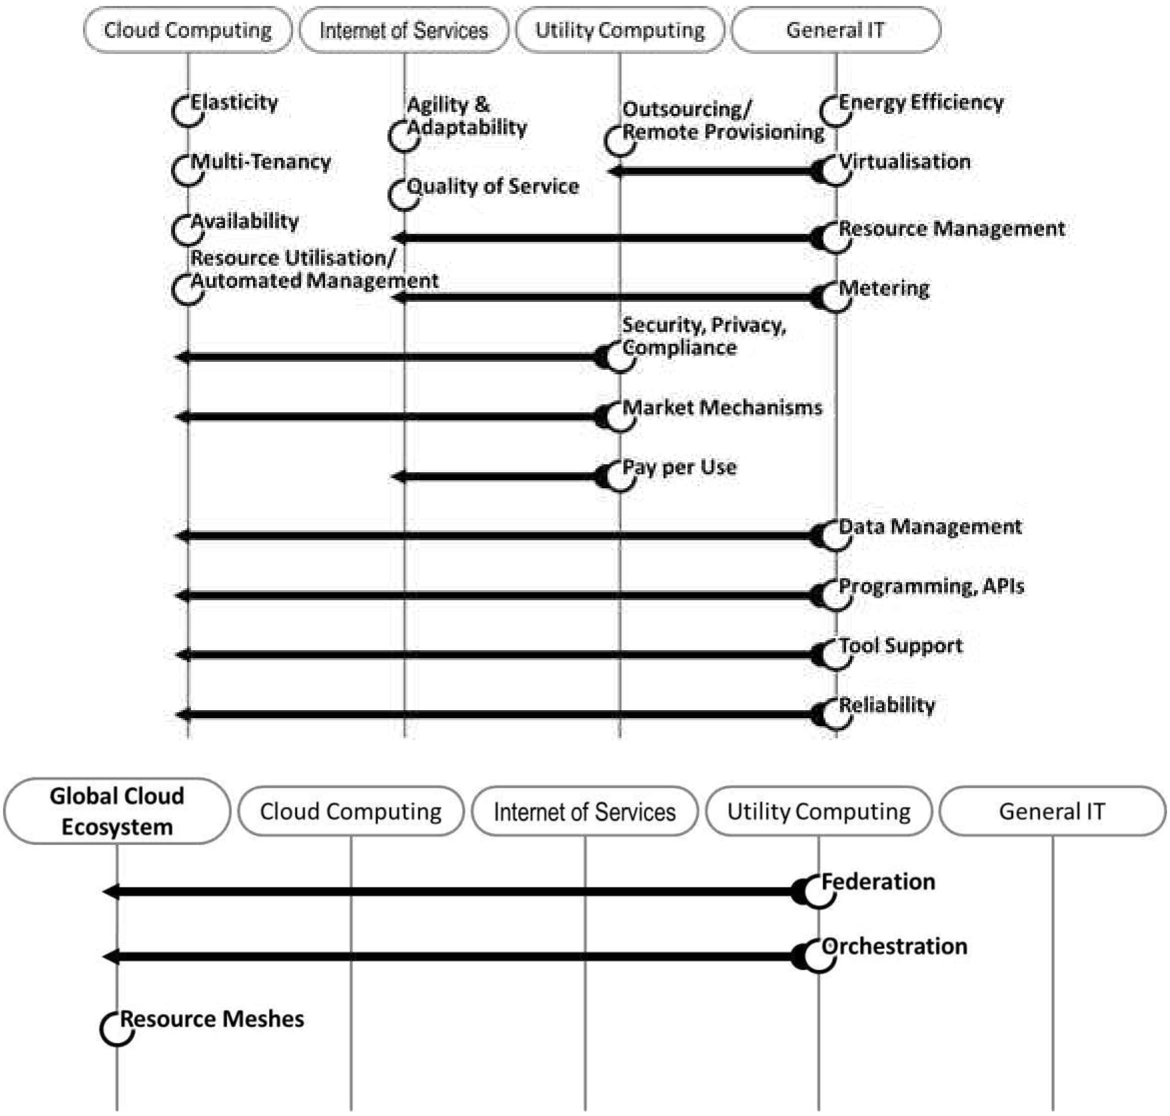

FIGURE 2: OVERVIEW OVER THE CLOUD CHARACTERISTICS AND THEIR RELATIONSHIP TO OTHER DOMAINS. INHERITANCE FROM RIGHT TO LEFT. THE FURTHER LEFT A CHARACTERISTIC ADVANCES, THE MORE IT NEEDS TO BE ADAPTED FOR THE RESPECTIVE DOMAIN(S). (TAKEN WITHOUT CHANGES FROM "ADVANCES IN CLOUDS", P.18 P.44)

It is therefore of utmost relevance that the multi-disciplinarity of, as well as the technical foundation behind CLOUDs is not only recognised, but actually exploited. This means in particular that the according disciplines should jointly address the outstanding research and development work, rather than repeating work already performed outside the own domain. It also means that the outstanding work can base on a solid foundation of research and development, that should be taken up and exploited. Furthermore, foundational advances can only be achieved with an according set of expertise behind it – expertise that is well-developed in Europe, yet tends towards similar fragmentation as the European industry.

C. TIMELINE

Europe has a great potential to not only participate in the CLOUD market, but also gain a long-term leading position in this sector. To this end, not only the specific pre-conditions of the European market, but also the progress made and the complexity of the outstanding research and development needs to be considered. Simply put: the technological status in Europe is not yet at a state which would allow realisation of the above scenarios right away.

In particular we can distinguish between work that should be immediately performed in close cooperation with the European industry in order to increase the competitiveness and capabilities of the current market players and encourage new entries into the CLOUD market. We classify this as "immediately relevant" actions that are required (a) to overcome the immediate barriers to participation in this market by European companies and (b) to get European companies involved in and thinking about CLOUD Computing. (a) is addressed particularly to the IT industry whereas (b) is generally applicable to all industry.

Whilst this type of actions provides a boost to industry, it helps only little in creating actual technical advances in all areas related to CLOUDs. Specifically, research work with a longer term perspective is required to provide and extend the necessary environment along with its tools and components to extend utilisation of CLOUDs in a usable, secure, cost-effective, performant and in particular sustainable manner. This type of work is primarily concerned with maintaining and increasing the market position gained through the "immediately relevant" actions. This is hence classified as "sustainability" actions.

Both types of R&D ("immediate" and "sustainable") are geared towards dealing with current technological obstacles and (entry) barriers, but are insufficient to address the substantial limitations of IT that will increasingly arise from the expected growth in IT usage, in particular over the CLOUD. Accordingly, research and development work more geared towards the long term development of IT is required to provide the foundations for paradigm shifts in several areas of technology and legislation where Europe could 'leap-frog' to an advantageous market position. All work of this type can be classified as "paradigm shifters" or "game changers".

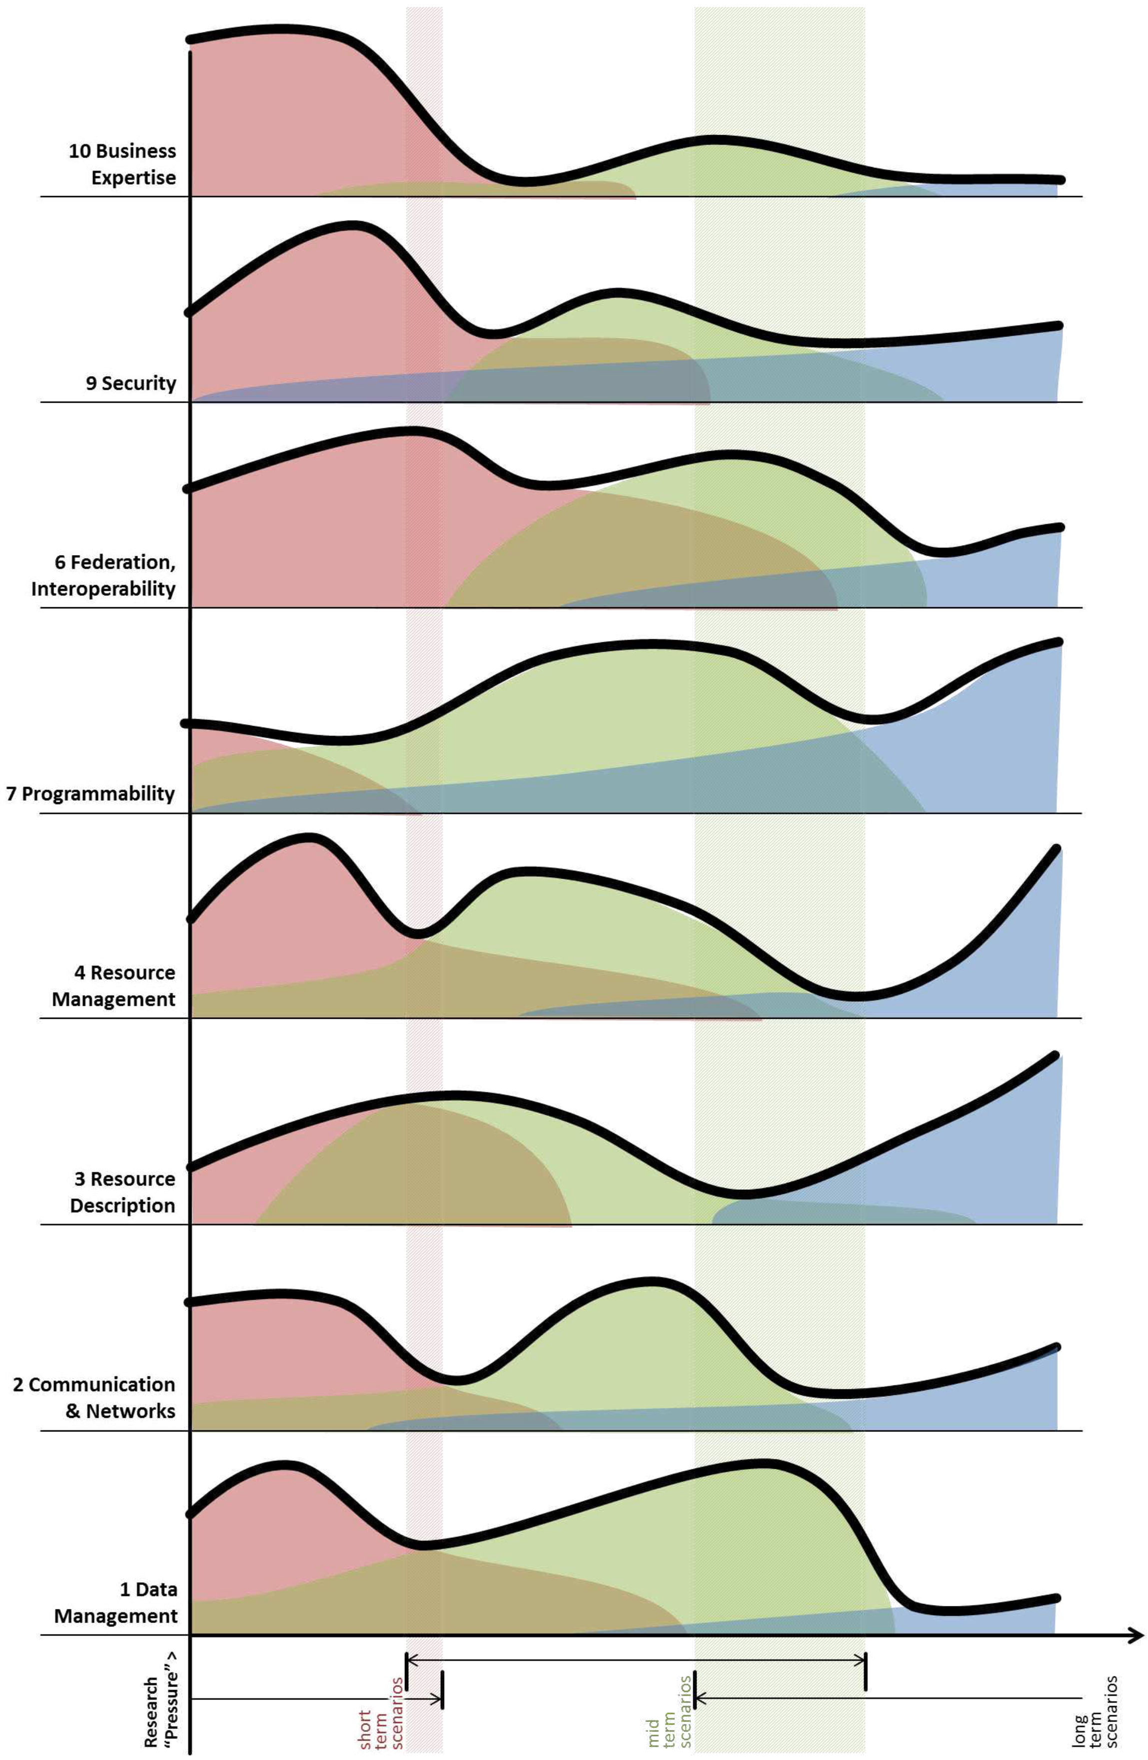

FIGURE 3: THE DEVELOPMENT OF R&D NEEDS OVER TIME

Even though these three types of R&D obviously become relevant in short-, mid- and long-term development of the IT industry, it must be noted that all types must be addressed at least partially right away, as they are each defined by different gestation periods, due to their complexity and technological pre-requisites.

It must thereby be kept in mind, that industrial development progresses and the market for CLOUD is subject to constant changes. Despite the turbulence in this development, however, success can be achieved by a clear roadmap that identifies strong long-term goals and a clearly structured path towards them. Such a roadmap needs constant updating, to adjust the details to reflect the changes in the market.

Figure 3 depicts the expected "pressure" put towards research due to the environmental development – in other words, it indicates the necessity to investigate into new / improved technologies in order to address the challenges posed by the scenarios discussed above. In most cases pressure builds up over time, until the scenario is expected to be reached. In the diagrams, this is depicted by the green, red and blue areas, with red being "immediately relevant", green "sustainability" and blue "paradigm shifters". It will be noted that frequently, pressure on research starts comparatively early due to the digestion time of the respective work – it is in these cases recommended to start with the according topics early to ensure availability at the time of need. Similarly, some research pressure will continue beyond the current scenario period, as the full scope of the topic may only become relevant at a later point.

Due to the nature of the assessment, it can be generally said that the accuracy of the prediction diminishes with the temporal distance. In particular with respect to the long-term paradigm changers, only a few concrete approaches can be outlined as yet Nonetheless, R&D on these topics now would place Europe in a leading position in 5-10 years' time:

1. Immediately Relevant

The following list of topics need to be supported so that the European IT industry can overcome the current barriers to entry to the market and so that European industry generally can utilise CLOUD computing for their business needs.

- Managing data deluge – volume, media types, streaming – to give confidence that CLOUDs can sustain appropriate data loadings and deal with the different media and streaming in an acceptable fashion. This requires in particular improvements in structuring data and mechanisms to handle big data;
- Improved networking through intelligent networking based on smart software to ensure provision of appropriate latencies within SLAs and to compensating off-line time inherent to the increased mobility.
- Improve the capabilities to exploit CLOUD features feasibly, which means in particular the elasticity behaviour of applications and a more integrate

monitoring-control loop between usage and application behaviour. This will demonstrate the advantages of having resources available on-demand whilst meeting the specific usage conditions.

- Similarly, improve the performance and portability of applications by understanding the resource types and their relationship to the use cases better. This can build up to a large extent on existing work in Grids and related domains.
- (Security) vulnerabilities need to be understood and handled better to overcome the lack of trust in utilising CLOUDs and the fears of security vulnerabilities and privacy invasions. This includes among others aspects such as improved identification and authentication, but also transparency for the user in layman terms.
- Reducing lock-in to proprietary solutions through means of increased portability and interoperability, to give organisations confidence that they can transport their application and data from one CLOUD infrastructure or platform provider to another as dictated by business economics and also can interoperate with applications hosted by other CLOUD providers (e.g. B2B).
- Support the competition, but also collaboration between CLOUD providers to give much greater flexibility in costing and pricing, performance, appropriate security etc. In the first instance this means comparability between offerings and their characteristics, as well as approaches to CLOUD bursting.
- Viable business models, respectively a clear European benefit is required so that the relevance of utilising CLOUDs is understood and the cost-benefit justified. To support generation of ROI, monitoring and accounting mechanisms need to be improved – this is a necessity for staying competitive in the market.

Proposed Instruments for Immediately Relevant R&D Work

Immediately relevant R&D topics aim specifically at boosting the current CLOUD industries, with a strong focus on SMEs. This means that these topics should be developed in close cooperation with industrial players. Any R&D work needs to clearly identify the industrial needs and provide concrete, directly usable solutions. Due to the different timelines of research and industrial need / development, any project in this domain should publish intermediary usable results in multiple iterations, to compensate as best as possible to the dynamic market situation. The project duration may implicitly be shorter than in other cases.

Potentially appropriate instruments:

- SME instrument H2020 (COSME continuation)
- Pre-commercial public procurement of technology
- Open calls under dedicated projects to generate quick innovation

2. Sustainability R&D

Some R&D is required to build up a European wide CLOUD ecosystem (cf. [SCH12]), in which services and data can be hosted and executed efficiently across heterogeneous CLOUD offerings. This allows users and providers not only to move their services and data freely between CLOUDs, but also to generate new types of capabilities through composition of various services, platforms and infrastructure. This gives European developers – especially SMEs – the advantage of a larger potential market (not locked-in to one platform) for their services and also for their services to e deployable through dynamic (re-)composition to provide optimal performance to meet SLAs.

To reach this objective, the following topics need to be addressed:

- In general, a better understanding of the relationship between application and user / usage behaviour, and the capability/ies to cater for this relationship efficiently. This will allow improved service provisioning, better resource utilisation and higher energy efficiency. This includes:
    - More enhanced metadata describing data, software, services, resources, users – to allow the CLOUDs middleware to manage service discovery, (re-)composition and execution autonomically;
    - Software engineering environment for CLOUD applications – to allow developers to produce quickly and cheaply well-designed and produced software for deployment in a CLOUDs environment;
    - Identification and provisioning of typical CLOUD usage patterns, in a wide variety of application domains. This also forms the basis for more complicated composed services as foreseen in the long-term vision.
- Enhanced mechanisms (based on rich metadata and related services) to deal with the growing heterogeneity within a CLOUD, as well as across multiple CLOUDs, thus paving the way for personalised service provisioning and usage.
- Support (through rich metadata and associates services) for extended outsourcing and composition of capabilities. Scalable orchestration of services and virtualised resources – to bring together the services required by an application with the required CLOUD resources at platform or infrastructure level to provide an appropriate execution environment; This must include mechanisms for enhanced protection and addressing privacy concerns to maintain trust in CLOUD provisioning.

Proposed Instruments for Sustainability R&D Work

Sustainability related R&D work should focus on ensuring that the industrial level CLOUDs achieved through the immediately relevant work can be maintained. This cannot be achieved simply through continuous short-term work, but requires a more forward looking approach that ensures the sustainability of the solutions, rather than the immediate solution to a problem. As part of the sustainability focus, all work in this domain needs to build up as much as possible on existing and accepted results. Such R&D work should integrate academia and industry equally: Industrial participation is needed in particular for specifying requirements and assessing usability / impact, as well as contributing to uptake. Academic research is needed for scientific advances, communication & education of next generation (sustainability of knowledge). The instrument is therefore similar to IPs and STREPs with considerable industry participation under FP7.

Potentially appropriate instruments:

- STREPs: focus on the improvement / realisation of a specific topic
- IPs: focus in particular on the multi-disciplinarity of CLOUDs

3. Paradigm Shifters & Game Changers

Once European industry has (a) gained a foothold in the market due to the short-term measures and (b) made further gains in quality and cost-benefit of service offerings through the mid-term (sustainability) initiatives, organisations should be ready to take advantage of the long-term R&D which provides new paradigms to address the essential long-term technical barriers and enable European industry to 'leap-frog' ahead in terms of capabilities, thus generating a massive advantage over other the competition.

The R&D work should enable European industry to easily provide (and exploit) a wide range of CLOUD use cases across a widely dispersed and highly heterogeneous environment, making best use of the full scope of resource types. This also means operating applications of (re-)composed services over federated heterogeneous (CLOUD-like) environments – including European intelligent mobile networking while the European IT industry supplies quality (software) services in a worldwide market.

- New paradigm for managing big data (integrity, state) – conventional data management preserves integrity and state. Heterogeneous distributed data provision and mobile code makes conventional techniques inapplicable.
- New paradigm for programming to be utilised by CLOUD middleware for optimal execution – conventional imperative programming styles with compilation and static type checking are less appropriate in a dynamic, heterogeneous, distributed CLOUDs environment. A new paradigm – probably declarative and with execute-time data binding – is required to cater for flexibility, heterogeneity, consistency etc.
- New techniques for interoperation and federation (heterogeneity and mobility) – the distributed heterogeneous nature of CLOUDs in the long-term time frame demands new methods for interoperation of applications (as services) across CLOUDs offerings and federation of multiple heterogeneous CLOUD platforms to appear to the application environment as one uniform platform.

- Exploitation of new device types and environments – advances in hardware for processors, storage, detectors and instruments and networks – along with associated low-level (embedded) software – will provide opportunities for exploitation for business benefit to which the CLOUDs environment must adapt in an agile, dynamic and flexible manner in order to give best cost-benefit.

Proposed Instruments for Paradigm Shifters

Paradigm Shifters / Game Changers are highly academic in nature: they need to explore approaches towards the long-term barriers, where a solution cannot be guaranteed. Industrial uptake and direct relevance are therefore more vague than for the other types. At the same time, R&D work of this type should maintain a high degree of "backward compatibility", i.e. with existing solutions, as much as possible to ensure smooth transition for industrial uptake. Industrial participation is still encouraged though, to ensure the principal value of the approaches pursued. Due to the nature of this type of research, and to allow for the "visionary" aspects of these topics, some failure must be expected and hence "allowed".

Potentially appropriate instruments:

- CSAs, Networks of Excellence
- "visionary" instruments similar to FET but with more "practical" focus and less risk involvement – large STREPs would be appropriate

D. NOTES ON INTERNATIONAL COLLABORATION

CLOUDs are a global phenomenon and not restricted to the European market or industry, or even academic environment: users and providers arise equally all over the world and want to make use of CLOUD based services without constraints of their country or other boundaries. As such, R&D must be tackled from the very beginning in a collaborative fashion with other countries, just for the reason of ensuring compatibility, interoperability and convertability across national and continental boundaries.

But also the actual market of service / resource provisioning and usage must consider cooperative approaches, rather than be seen as a purely competitive environment, similar to the joint efforts in roaming - only this way the user demands can be fully satisfied in the first instance. And also similar to roaming, this process of generating a global CLOUD ecosystem will be slow, obstructed by multiple barriers including a lack of common standards, policies, legislative compatibility, sensible common (respectively overlapping) business models and comparability of the offerings.

At the same time, however, the market will for some time continue to grow in a competitive fashion, which necessitates early engagement in joint efforts, but, more importantly, also close observation of activities in non-European countries, lest their advances should become world-dominating.

In general, international collaboration on all R&D levels is encouraged not only so as to meet the global vision of CLOUDs, but also for the simple reason that no country in the world can claim scientific and technological dominance over any domain, let alone over all IT related disciplines. Joint efforts are mandatory simply in order to make the necessary technological advances happen and break down the associated barriers effectively and in a sustainable fashion. Any approach dominated by a single domain must implicitly face the problem of having to redevelop expertise already existing for multiple decades.

We must in this context distinguish between the industrial and the academic specialisation of individual countries. Whereas the first can be seen as an indicator for the specific strengths & opportunities of a specific country in the global market, and thus for identification of the appropriate market niches in a global ecosystem, the second must be regarded as an opportunity for the exchange of knowledge and expertise.

We can distinguish industrial and academic dominance in different regions / countries in the world. For example, it can be noted that there is currently a stronger hardware market in the USA, than in Europe, whereas Asia is strong in end-user electronic appliances [ROS11]. Europe on the other hand has a strong telecommunication industry, and a general dominance in most engineering (see e.g. [ECO12]).On the other hand, Europe has a strong academic background in distributed computing and automation, whereas Japan for example is leading in robotics.

Obviously, there are multiple domains where no clear dominance can be assigned to any country, unless some distinction is made on a very fine granularity.

In particular countries that currently exhibit a strong industrial growth or such a potential, yet are still comparatively weak in the global economy, offer a considerable potential for joint efforts to build up a stronger competitive backbone for Europe's (resource) infrastructure. The apparent American dominance in many CLOUD building blocks creates an environment where other continents may need to consider close cooperation for balancing the ecosystem. SMEs can thus principally build up and serve a customer base, they would not reach with their own resources alone. This relates to countries such as China, Russia, Brazil and Mexico, but also Africa[8] and India in particular in the mobile industry. China for example has demonstrated during the past decades strong and steady growth both in economy, science, and technology. Considering the great importance of China to European economy, and its dependency upon European experience and expertise, it is evident that there are

---

8 E.g. http://www.accenture.com/SiteCollectionDocuments/Local_South_Africa/PDF/Accenture-The-Dynamic-African-Consumer-Market-Exploring-Growth-Opportunities-in-Sub-Saharan-Africa.pdf

many opportunities for mutually beneficial cooperation. Such joint industrial effort relates specifically to R&D actions of the “immediate” and “sustainability” type.

On the other hand, merged efforts from the academic side can help particularly to address the common challenges arising from long-term changes in IT, i.e. for addressing the necessary paradigm shift. Common research objectives should thereby in particular focus on either (a) exploiting academic expertise beyond European boundaries or (b) on concrete topics that require collaboration to realise the federated long-term vision, such as common standards, legislative aspects etc.

Thus, some of the specific topics that may be explicitly interesting per country include for example federated services offerings with Brazil, logistics with Japan, interoperability concerns with Brazil and the United States of America, education and eLearning together with India etc. Obviously, many further opportunities exist in which collaboration may be fruitful for all parties involved – these have to be investigated on a case by case basis.

III. CONCLUDING REMARKS

The current trough of disillusionment in CLOUDs is a major opportunity for Europe to "leap ahead" when other countries struggle to define their concrete position in CLOUDs. A casual observer might assume the CLOUD problem is solved because there are commercial solutions from across the Atlantic. However, a closer look indicates that these solutions are proprietary, incompatible, not necessarily optimal for European business and not sustainable in the long run given the growing technological issues that arise from all the changes in IT and its usage.

Europe's heterogeneity in infrastructures and industrial players, in particular SMEs, plays a driving force and major potential in this context. Thanks to the strong telecommunication industry, Europe also has to struggle less with unifying the underlying layer and thus can develop solutions faster. With this telecommunication industry, Europe also has a strong leverage on competing with the currently dominant CLOUD infrastructures.

The long-term vision and relevance of CLOUDs is neither constrained to a specific country, nor a specific technological domain. To fully exploit the potential of CLOUDs, it is therefore mandatory to work hand in hand with other countries and ensure that approaches are interoperable across (legal, administrative and political) boundaries. Commonalities between countries can be exploited to increase the market competitiveness and value. Similarly, CLOUDs are not defined by a single technological domain. Due to their nature, they incorporate multiple technological aspects. Solutions in CLOUDs are therefore not only valuable for the Software industry in Europe, but for a wide area of IT technologies and users. The CLOUD Vision can thus only be realised if policy makers, technical analysts and industrial stakeholders work together.

The value of more forward looking R&D for the European industry cannot be overrated. Whereas Europe needs to (a) do some short term commercially relevant developments to catch up, and (b) do some mid-term R&D for sustainability to keep up and also to initiate some advanced features of CLOUDs needed by Europe especially in platform provision, interoperation and service development environments, Europe also needs to (c) initiate some longer-term research leading to superior platforms, interoperable platforms and novel services to intercept the (by then sophisticated) market demand in the 10-year horizon. **All three threads of activity must start very soon**.

REFERENCES

[ECO12] ECORYS (2012). EU SMEs in 2012: at the crossroads. Annual report on small and medium-sized enterprises in the EU, 2011/12. Available online at http://ec.europa.eu/enterprise/policies/sme/facts-figures-analysis/performance-review/files/supporting-documents/2012/annual-report_en.pdf

[FEN11] Fenn, J. & LeHong, H. (2011). Hype Cycle for Emerging Technologies 2011, Gartner. Available online: http://www.gartner.com/DisplayDocument?doc_cd=215650

[KRO12] Neelie Kroes, Vice-President of the European Commission responsible for the Digital Agenda Setting up the European CLOUD Partnership World Economic Forum Davos, Switzerland, 26th January 2012. http://europa.eu/rapid/press-release_SPEECH-12-38_en.htm

[ROS11] Rospide, S. (2011). Electronic Industry Outlook 2010 – 2015 - Smart cities, from concepts to markets. Available online at: http://www.decision.eu/doc/presentations/Decision_Rospide_Summit_2011.pdf

[SCH10] Schubert, L., Jeffery, K. & Neidecker-Lutz, B. (eds). (2010). The Future of CLOUD Computing – Report from the first CLOUD Computing Expert Working Group Meeting. Cordis (Online), BE: European Commission. Available online: http://cordis.europa.eu/fp7/ict/ssai/docs/CLOUD-report-final.pdf

[SCH12] L. Schubert, K. Jeffery (eds.) (2012). Advances in CLOUDs - Research in Future CLOUD Computing, Cordis (Online), BE. European Commission. Available at: http://cordis.europa.eu/fp7/ict/ssai/docs/future-cc-2may-finalreport-experts.pdf

[WSH12] Communications Networks, Content and Technology Directorate-General Unit CONNECT A3 – Complex Systems and Advanced Computing (2012). Towards a Breakthrough in Software for Advanced Computing Systems - Report from a Workshop organised by the European Commission in preparation for HORIZON 2020. Available online at: http://cordis.europa.eu/fp7/ict/computing/documents/advanced_computing_ws_report.pdf

European Commission

A Roadmap for Advanced Cloud Technologies under H2020

Luxembourg: Publications Office of the European Union

2012 — 21 x 29,7 cm