Report on the public consultation for H2020 Work Programme 2016-17: Cloud Computing and Software

Web-based consultation: 10 September-17 October 2014

Workshop date: 4 November 2014

Final version

4 December 2014

Rapporteur: David Griffin

SECTION 1. EXECUTIVE SUMMARY

The Software & Services, Cloud unit (DG CONNECT E.2) undertook an open web-based public consultation on the research and innovation challenges for the Horizon 2020 Work Programme 2016-17. Sixty-one written contributions on Cloud Computing and twenty-six on Software (including Open Source) were submitted during the consultation period. A post-consultation workshop was held on 4 November 2014 where the contributions were presented and discussed to identify research priorities.

The consultation contributions and the workshop discussions highlighted significant current research challenges and priorities in both Cloud Computing and Software (including Open Source), as summarised in the following paragraphs.

Research priorities in **Cloud Computing** cover several major trends such as: cloud interoperability and federation; security, privacy and confidentiality for cloud-based systems; and new engineering methodologies for developing/deploying cloud-based software.

Extensions to cloud computing were identified including: exploitation of specialised hardware components; the management and control of network resources within and between data centres; rolling out the cloud to include edge computing; and the inclusion of sensors and actuators to support applications and services for the Internet of Things.

Attention should be placed on the non-functional aspects of cloud-based applications and services including SLAs and performance optimisation of real-time, cloud-based applications and services, and increased energy efficiency through the exploitation of hardware advances, lightweight virtualisation techniques, and intelligent auto-scaling algorithms.

Application-specific challenges identified by the consultation process include mechanisms for automated service assembly; support for big data & high performance computing applications; and services for vertical markets.

Advanced platforms to experiment with cloud services and deployment configurations in realistic environments were seen as a key tool for future projects as were simulation techniques. They will enable experimenting at scale with application deployment in highly distributed and federated clouds.

Several policy and regulatory topics in cloud computing were raised related to the need for EU-based cloud providers and a coherent legal framework on the security and privacy of cloud computing.

The challenges and research priorities in **Software (including Open Source)** covered all phases of the software development cycle, starting with new approaches, languages and tools for requirements engineering; agile/scrum development frameworks; software development methodologies and tools for IoT and other highly distributed systems; higher-level, intention-driven programming models; and approaches for the automation of interoperability & compliance testing, risk identification, detection and mitigation.

New tools and mechanisms are required for managing software quality including the use of big data analytics on information from a variety of sources including user feedback data and software performance monitoring metrics.

Finally, the consultation identified the need for a wider adoption of open standards and the need for reliable and sustainable open source repositories with appropriate governance mechanisms to aid exploitation actions of EU projects producing open source software.

TABLE OF CONTENTS

SECTION 2. OUTCOME OF THE CONSULTATION

This section of the report summarises the research and innovation priorities identified by the written contributions and raised by the discussions during the post-consultation workshop.

The first part describes the challenges in Cloud Computing while the second focuses on Software (including Open Source).

A more detailed analysis of the contributions and workshop discussions leading to the identified challenges is contained in section 3, on Cloud Computing, and section 4, on Software (including Open Source).

2.1 Cloud Computing

Despite **cloud federation and interworking** being the subject of current research projects, many issues remain open, including: techniques for load balancing, off-loading and migration between cloud systems; resilience of services deployed over federated clouds; automated rule-based security and privacy enforcement and traceability in multi-actor environments; cloud brokers; and mechanisms to avoid vendor lock-in and for switching between cloud providers. Mechanisms for improved data portability and interoperability are required across public, private and hybrid cloud infrastructures.

Security, privacy and confidentiality issues continue to be seen as high priority research themes to increase trust in cloud computing and foster the wider uptake of cloud-based applications and services, including: verifiable and measurable chains of trust and accountability; intrusion detection systems within virtualised environments; application and data security in multi-tenant environments; greater user control of how data is stored, routed and accessed within clouds and across domains, including controls for deleting confidential data; improved techniques for mapping high-level security/privacy requirements to underlying technical solutions (IaaS, PaaS, data management); development of techniques such as blind storage, homomorphic and end-to-end encryption; mechanisms for managing personal identity in digital space across heterogeneous systems and across multiple cloud providers.

Engineering methodologies need to be advanced for developing/deploying cloud-based software, including: intention-driven software engineering methods; support for large-scale data processing including security analytics for identifying unexpected system behaviour; resource usage modelling and optimisation; approaches to improve the reliability of high performance RAMCloud-based data centres; new agile dynamic system engineering methods; techniques to facilitate software migration across heterogeneous infrastructures; techniques to support the co-evolution of database architectures and cloud computing; continuous development approaches for cloud-based software inspired by DevOps.

Further work is required on expanding the management and configuration of **specialised hardware processing components** such as GPGPUs and general-purpose ASICs in cloud infrastructures. Advanced control and **management of networking resources** within and between data centres is required to improve the efficiency, scalability and resilience of distributed cloud systems, especially in the context of ultra-low latency and high capacity network technologies.

New approaches are required to integrate cloud computing with cyber-physical systems and mobile devices to create a **more decentralised cloud infrastructure** extending the **edge/fog computing** paradigm to the extreme edge of the network.

Integration of sensing and actuation resources with cloud computing is required to unlock new value-added **cloud-based IoT services**. Further research is needed into lightweight techniques for IaaS deployment on constrained computing platforms and on PaaS functionality for stream and event processing for IoT applications.

Research should be conducted into how **more stringent SLAs** can be supported in clouds to support **mission-critical and high-performance applications**, and on finer-grained control of resources to optimise the performance of real-time applications for cyber-physical systems, IoT and safety-critical systems.

Energy efficiency in cloud computing can be improved through the exploitation of hardware and firmware advances, by adopting lightweight virtualisation, and through the use of intelligent auto-scaling algorithms for adapting cloud-based software deployments to use minimal resources while complying with SLAs.

More research is needed on **self-managing and autonomic services** running within the cloud and for evolving DevOps environments for managing cloud-based systems.

Programming languages and engineering frameworks need further development to support applications that can migrate to any device and exploit the parallelism of **distributed, heterogeneous clouds**. Tools for automated, contextual instantiation, deployment, configuration and life-cycle management on multiple heterogeneous cloud platforms are needed.

Tools and techniques are required for the design of **big data and HPC applications** including: models for a better understanding of specific types and large quantities of data; data compression and transmission techniques; self-adaptation mechanisms for optimising the performance and efficiency of large-scale data processing applications in the cloud.

Platforms are required to experiment with cloud services and deployment configurations in realistic environments for both research and business objectives. Flexible platforms will enable experimenting at scale with decentralised and federated approaches, investigating the performance of advanced hardware and firmware features. Integration of EU-based experimental platforms with other international initiatives will allow global experimentation at large scale. Further discussion is needed to identify longer-term features and technologies to be supported by platforms for experimentation, beyond those supported by existing infrastructures in the EU and those available in the USA/Asia.

Tools are needed for simulating, during design and development phases, the behaviour of applications under a variety of deployment configurations in the cloud. **Simulation tools** are also required to assist with the development of dynamic reconfiguration strategies and scaling algorithms.

Open standardisation of service, platform and infrastructure interfaces is required to avoid the inefficiencies caused by a plethora of technology-specific and platform-specific APIs.

Measures to increase the uptake of cloud computing include: appropriate abstractions of the infrastructure and computational complexities of the cloud; automation of cloud-based software deployment and lifecycle management; new tools and programming models to assist with software migration from multi-core systems to the cloud.

Novel collaborative business models will facilitate interactions between stakeholders in federated and distributed clouds allowing individual SMEs to concentrate on their specialisation in the business chain and aid European SMEs in gaining a global market reach for their technologies and products.

Several contributions identified a dependency on large cloud providers from outside of Europe. Establishing effective **federations of public cloud providers within the EU** would allow data (and revenue) to remain within the EU and be subject to European standards, certification and legislation.

Several contributions referred to **regulation and EU policy** aspects of cloud computing, calling for a legal framework on the security and privacy of cloud-based data.

2.2 Software (including Open Source)

Further attention should be put on **requirements engineering**. Use of natural language for requirements capture is inherently ambiguous and unstructured: research into new methods of requirements capture linked with CASE systems is required. New techniques and tools are required for compiling requirements to final systems, abstracting underlying technological platforms.

Higher-level, **intention-driven programming models** need to be developed for software replacement, recomposition and adaptability.

New tools are required to support modern **agile/scrum development frameworks** to assist with developing mission-critical and safety-critical systems. Further studies are required on the software monitoring systems and metrics applicable to agile/scrum methods.

Further research is required into tools and mechanisms for specifying, assessing and **managing software quality**; to make use of contextual data collected at runtime through various sensors, online data sources, smart objects/appliances, developer forums and social networks.

Big data analytics of user feedback data, software performance monitoring and data from other sources can be used to gain insight on user trends and preferences, measure application performance and resource consumption metrics, identify feature and performance improvement opportunities and the causes of software failures.

Development environments are required which treat the cloud as a single virtual machine enabling scalable applications or services running in distributed and heterogeneous environments to be developed as easily as if they were running on a single-processor computer. New algorithmic and processing models for distributed, heterogeneous execution environments are also required, for example supporting processing without immediate access to data. Tools should also facilitate the porting of legacy software to the cloud, addressing issues such as remote user interaction, security, and data transfer.

Model-driven software engineering tools are required to aid the automation of interoperability and compliance testing for cloud-based applications and services. New scalable approaches are needed for information system design and analysis based on automated risk identification and analysis, detection and mitigation.

Despite the proliferation of semantic web services more work is required on techniques for **automatic composition of services** by intelligent software agents to achieve a targeted task.

Reliable and sustainable **open source repositories with adequate governance mechanisms** are required to aid exploitation actions of EU projects producing open source software. These actions should be supported by further research on the provision of quality testing platforms for open source software.

Contributions to **open software standards** should be undertaken to speed up standardisation of APIs and linked open data.

New methods of developing applications under **weak data consistency** are required for stand-alone multi-core systems and for distributed cloud environments and studies are required on the performance and efficiency trade-offs of supporting weak data consistency at the application layer versus being natively supported by the infrastructure.

New frameworks are needed for **software development for IoT and other highly distributed systems**. Trust was identified as one of the issues to be addressed but also the need for interfaces to discover where data is stored/collected and its legal context.

SECTION 3. DETAILED ANALYSIS OF RESEARCH CHALLENGES AND PRIORITIES IN CLOUD COMPUTING

The following sub-sections present the key research challenges and priorities for the H2020 ICT Work Programme 2016-17 on Cloud Computing as proposed in the written contributions and discussed in the workshop. The sub-sections are organised as follows:

- First of all, the major trends in cloud computing are addressed in section 3.1 on cloud interoperability and federation, section 3.2 on security, privacy and trust, and section 3.3 on software engineering for the cloud.
- Secondly, extensions of cloud computing to encompass underlying hardware and networks are addressed in section 3.4, rolling out the cloud to the edge in section 3.5 and encompassing the Internet of Things in section 3.6.
- Sections 3.7 and 3.8 cover non-functional aspects of performance and energy efficiency;
- Section 3.9 identifies the challenges for the management of clouds and cloud-based applications.
- Sections 3.10, 3.11 and 3.12 discuss cloud applications from the point of view of service assembly, big data & high performance computing, and vertical markets.
- Section 3.13 discusses experimentation for cloud-based applications and services while section 3.14 addresses simulations to support application development and deployment.
- Finally, policy-related aspects are covered in section 3.15 on standards, section 3.16 on increasing the uptake of cloud computing, section 3.17 on business models, section 3.18 on European clouds, and section 3.19 on regulatory and policy issues.

3.1 Cloud Interoperability, Federation and Brokerage

Despite cloud federation and interworking being the subject of current research projects, many issues remain open and require additional research and innovation activities [C-8] [C-19] [C-58]. Ongoing research issues include: automatic and efficient cloud management, including load balancing, off-loading and migration patterns for computing and data between cloud systems [C-58]; application and service resilience in federated cloud deployments [C-56]; automatic-rule based security and privacy enforcement and associated traceability for data and computation in multi-actor environments [C-58]; seamless integration between public and private clouds through convergence of underlying cloud technologies, management tools and methodologies [C-8] [C-49] [C-58]; actions towards reducing the gap between characteristics of public and private clouds [C-10] [C-58], enabling hybrid private-public cloud solutions [C-28] as well as enabling multiple overlapping community clouds [C-50] [C-56]; cloud brokers supporting cloud marketplaces are required to assist users with selecting the right service for their needs, benchmark/compare cloud provider capabilities & performance levels, and consolidate billing, thereby avoiding complex sets of supplier-user relationships [C-10] [C-12] [C-46] [C-51]. In addition, broker-less models should be developed for federation, interworking and dynamic provisioning over multiple cloud providers [C-51]. Mechanisms for automated and secure load sharing across federated data centres with tools for performance identification/monitoring are also required [C-46].

Accessing and reasoning with data and functionality from different systems typically requires a significant amount of manual intervention, which does not easily scale to multiple clouds, so methods for reducing manual effort for interoperation are needed. Further work is required on standard interfaces and APIs for describing resources, accessing, configuring and monitoring cloud functionality [C-28] [C-47]; integrated data management systems for multi-site clouds with heterogeneous data stores [C-48]. Furthermore [C-7] proposed the use of machine-interpretable mark-up to describe interfaces/APIs and underlying cloud function semantics to enable customers to switch between cloud providers and substitute one resource for another, helping foster a pan-European competitive market for cloud products. Pragmatic approaches in this direction include the management of a catalogue of trusted cloud provider capabilities, technologies and standards [C-50] [C-56].

Mechanisms are required to guarantee data portability and interoperability across different public, private and hybrid cloud infrastructures [C-18] [C-46] [C-47] and between cloud-based services/applications and heterogeneous mobile devices [C-46] [C-54]. Solutions addressing security and integrity issues are required when data is transferred between cloud-based applications and when using different APIs to access data across multiple cloud providers [C-8] [C-22]. Furthermore, approaches and agreements are required for the long-term preservation of data [C-47] especially in the case of migration between providers. Users should remain the owners of their data stored in cloud storage systems [C-9].

Several contributions identified the need to avoid vendor lock-in [C-19] [C-28] [C-50] and to ensure properties such as interoperability, service level agreements and security when cloud services are operated across multiple domains [C-9] [C-56]. Barriers to customers running applications and services over multiple cloud solutions are due to a lack of common architectural patterns, open standards and APIs for interoperability & portability, as well as security and privacy concerns. Developing cloud-based applications with components that can be dynamically adapted and deployed on multiple cloud providers at a time, taking into consideration requirements such as cost, security, regulations and QoS, presents a significant number of challenges to be overcome [C-33].

Several contributions identified prior experience of designing and deploying federated cloud systems [C-48] [C-50], including a federation of private clouds across Europe by applying its expertise in operating a federated scientific Grid infrastructure [C-47].

3.2 Security, Privacy, Identity and Trust in Cloud Computing

Security, privacy and confidentiality issues continue to be seen as high priority research themes to increase trust in cloud computing and foster the wider uptake of cloud-based applications and services [C-6]. Security and privacy should not be the concern of users but embedded in cloud systems by design [C-52]. Techniques to increase trust include: establishment of a catalogue summarising the available trusted cloud services across Europe, operated on a not-for-profit basis managed by an independent entity working to create a new channel for market providers [C-56] [C-61]; establishment of a trust chain of accountability for data in the cloud and a widely understood and verifiable trustmark “nutritional label” or measurable security properties that presents to consumers the ways organisations collect, use, and share personal information [C-24] [C-61]; mechanisms to assure that shared data is complete, valid, consistent, accurate, precise, and up-to-date [C-21]; intrusion detection systems within virtualised environments using self-adaptation techniques while preserving performance levels [C-48], with attention to application and data security in multi-tenant environments [C-22].

Further research is required on techniques for greater control by users of data traceability, allowing individuals control of how their data will be stored and routed, with mechanisms for notifications if data are transferred outside specified boundaries [C-25] [C-26] [C-38], or have been accessed or altered without authorisation [C-34]. Particular emphasis should be put on sensitive information such as medical data [C-19], but, on the other hand, this should not hinder advances in high performance computing and big data analytics for important fields such as personalised medicine or genomics that require mobilisation of large quantities of resources in federated clouds [C-44]. Verifiable mechanisms are needed for greater transparency on how cloud providers manage customers' personal data and/or copyrighted data on a cloud system [C-38] [C-41] [C-46] [C-56]. Future research should aim at addressing the governance of data and provide stakeholders with the means to exercise appropriate controls on their data so that they may access, validate or delete data when needed [C-56].

Improved tools and techniques are required for the mapping of high-level security and privacy requirements (c.f. Trusted Cloud Europe) to the underlying deployed technical solutions (IaaS, PaaS, data management) [C-58]. Advanced techniques such as blind storage, homomorphic encryption and end-to-end encryption at the network level are required to be studied and deployed for security and privacy control on data and computations [C-15] [C-16] [C-17] [C-58] together with techniques for privacy-preserving cryptology including anonymous credentials and practical techniques for processing encrypted data. Furthermore, research into functional encryption such as attribute-based cryptography and cryptography with outsourced services should be emphasised [C-56]. Encryption keys and encrypted data should be visible only to computation nodes that are trusted by the final user [C-38].

Currently there is no single way to precisely determine the identity of a person or entity in digital space. Contributions proposed work on establishing a European Digital ID and a verified European system of identification [C-55]. This would facilitate common access control and identity management mechanism across all cloud infrastructures, including federated cloud platforms [C-47] [C-50]. Financial transactions would also benefit from a common digital ID [S-25]. Security and admission control with inherent support for the identification of individuals, devices and organisations, resource access control is required as clouds are increasingly interconnected and are integrated with heterogeneous systems and devices [C-58]. The workshop discussions highlighted that a lack of identity management between collaborating clouds is hindering business between EU countries and that reliable and easy-to-use systems for building, managing and using federated trust networks across multiple clouds are required.

The post-consultation workshop identified that security and confidentiality in cloud computing could be increased by users encrypting personal and corporate data before it is uploaded to cloud infrastructures, rather than private keys needing to be held centrally by cloud-based cryptographic services. This will require enhanced security awareness by private and corporate users of cloud services.

3.3 Software Engineering for Cloud-based Systems

This subsection addresses the identified software engineering challenges for cloud-based systems as submitted to the consultation on Cloud Computing. There is a related set of challenges in section 4.2 on Software Development for the Cloud which reports on the contributions to the consultation on Software (including Open Source).

Engineering methodologies need to be advanced for developing/deploying software via the cloud. Intention-driven software engineering and description models should be extended to include non-functional attributes such as portability [C-22]. Advances in cloud-based software engineering technology are required to support very large-scale applications such as massive parallel image processing and data analysis [C-19]. New models are required for describing cloud applications covering a description of the structure of an application and the description of its expected behaviour to allow cloud providers to better optimise their resource usage [C-

48]. New approaches are needed for RAMCloud job scheduling on data centres with no hard disk storage and are therefore susceptible to data loss [C-46].

New agile dynamic system engineering methods are needed for the definition, composition and deployment of complex cloud-based applications and services. Conventional information system design and analysis approaches are not well suited to managing risks in dynamic and evolving systems. Approaches based on expert analysis cannot be scaled and sped up enough to be used in large-scale, autonomically provisioned and managed cloud infrastructure and applications. New approaches are needed for incorporating models, methods and tools in a coherent fashion throughout the life cycle of cloud systems and applications, from early design to final decommissioning. A coherent approach should support automated risk identification and analysis, detection and mitigation that can be applied at scale during design and operation of cloud-based virtualised infrastructure and services [C-20].

Novel approaches are required for programming multi-core and distributed cloud-based systems on weak data consistency. Ideally, multi-core and cloud-based data processing and storage systems should provide strong consistency, meaning that replication is transparent and the system behaves as though it was processing client requests serially on a unique copy of the data. Strong data consistency provides an intuitive programming model for application developers, but achieving it requires replicas to synchronise, negating the benefits provided by distribution. Writing correct applications under weak data consistency is challenging and new methods are required for: formulating precise specifications of the guarantees provided by weakly consistent systems; developing automated analysis tools for finding bugs and verifying critical properties of deployed systems; developing software productivity tools for weakly consistent cloud storage. Low-level solutions for weak data consistency in multi-core systems needs further study and the way weak data consistency could be supported in distributed cloud environments remains a research issue [C-30] [S-16].

New programming models leveraging cloud infrastructures are required to cover: data flow, event processing, big and fast data processing, highly distributed or parallel systems; "cloud proof" languages with clear semantics for distribution, failures, parallelism, event processing; impedance matching between programming languages and the APIs used to program and manage the cloud; including abstractions for cloud architectures and behaviours, consistently addressing the full XaaS stack to reduce the vertical heterogeneity of clouds; optimisation and transformation techniques to facilitate software migration across heterogeneous infrastructures; covering the co-evolution of database architectures and cloud computing to address new challenges on scale, resilience, privacy and security with advanced optimisation for cloud-data coordination and in-memory data processing [C-54] [C-58].

New approaches are needed for engineering socio-technical cloud-based software, such as systems for healthcare, emergency and disaster management, sustainability, smart cities, supporting organisational work, e-commerce and e-business collaboration and teamwork and social networks. Issues to be addressed include how to express accountability relationships formally: recognising accountability relationships as first-class software abstractions is the first and essential step toward systematically building socio-technical systems. [C-27]

Adoption of continuous development approaches for cloud-based software inspired by DevOps [C-5] [C-8]. Typically, legacy applications were not designed with multi-tenancy in mind. Porting this code to a SaaS environment might expose problems with data isolation. Additionally, newly developed or modified applications, especially those developed under the time pressure of short development cycles, may accidentally contain security holes. Even worse, applications made either freely available, or for a low cost, such as in a mobile marketplace, may have been maliciously developed with back-door security breaches. Big data security analytics for identifying unexpected system behaviour holds the promise to overcome many of the shortcomings exhibited by earlier generations of system protection, such as signature-based detection [C-5].

Evolutions of the IaaS/PaaS/SaaS model are required to support highly dynamic ecosystems of increasing complexity, mixing layers (IaaS, SaaS...) and functionalities (compute, data, networking) [C-58]. Discussions in the post-consultation workshop identified that lean/lightweight programming platforms are required for cloud-based software systems that cover multiple layers from the network layer upwards.

3.4 Hardware and Network Support for Cloud-based Applications and Services

Further work is required on expanding the portfolio of specialised hardware processing components such as GPGPUs and general-purpose ASICs that can be virtualised in cloud infrastructures [C-47] [C-48], including algorithms and techniques for resource allocation and configuration of the lower building blocks of the cloud (such as VMs, IO devices, memory and multi-core, multi-machine scheduling) to improve the management of physical resources and impact the efficiency of cloud systems in terms of energy and costs [C-48] [C-53] [C-58]. Also advanced control and management of networking resources (including network virtualisation, SDN and NaaS interfaces) within and between data centres is required to improve the efficiency, scalability and resilience of distributed cloud systems [C-58].

PaaS orchestration and management systems should develop new models for the optimisation of virtual network resources, making use of SDN and novel publish-subscribe mechanisms for inter-data centre operations, for optimising service placement, scaling and load balancing capacities, including interoperability in federated PaaS environments [C-52].

3.5 Edge/Fog Computing and the Mobile Cloud

New approaches are required to integrate cloud computing with cyber-physical systems and mobile devices to create a more decentralised cloud infrastructure extending the edge/fog computing paradigm to the extreme edge of the network. Cheap, low-power consumption nano-PCs deployed at home as gateways to cloud services to improve performance for low-latency applications, increase data privacy, reduce data-access fees and improve energy efficiency [C-48]. Mechanisms are required to be developed to integrate non-traditional edge systems and devices as part of the cloud environment so they can be treated as extended cloud resources in heterogeneous ecosystems. Tools and techniques are needed to design, develop and deploy software applications in such heterogeneous, highly distributed environments while taking advantage of cloud characteristics such as elasticity, scalability, distribution and automation [C-8] [C-33] [C-46] [C-58].

3.6 Integration of Cloud Computing with the Internet of Things

Merging cloud computing with IoT to deliver new value-added services requires further work on the integration of sensing and actuation resources as abstracted, virtualised and grouped resources [C-22] [C-52] [C-58]. R&D activities are needed in order to develop a comprehensive framework implementing tools for smart devices integration, configuration and management [C-11], with widely distributed and heterogeneous data sources, extending cloud computing towards connected devices and systems in an effort to provide consistent and managed environments for IoT services [C-54] [C-56]. Furthermore, IaaS deployment is required on constrained computing platforms with dedicated PaaS functionality including data, stream and event processing for IoT applications [C-58]. Legally and globally consistent identifiers for smart devices, sensors and actuators are required [C-46].

3.7 Predictable Performance for Cloud-based Systems

Although there has been much academic discussion of SLAs for cloud services, cloud providers have only really offered very simple SLAs with a take-it-or-leave it attitude. Research should be conducted into how more stringent SLAs can be developed on top of existing weaker SLAs [C-9] [C-32]. Consistent SLA support for mission-critical applications is required at provisioning and run time with a continuous optimisation of resources [C-58]. Each application in a cloud infrastructure receives a fraction of the total computation power and current providers cannot provide a strict and formally proved isolation among concurrent applications – further work on the fine-grained control of underlying cloud infrastructure and resources is required to optimise the performance of real-time applications for the integration of next generation of cyber-physical systems, IoT and safety-critical systems. A probabilistic approach would be able to trade-off the level of guarantees with the amount of resource over-provisioning to improve the predictability of cloud-based applications without losing the cost benefits of the cloud [C-60].

Important cloud-based applications/services for stringent SLAs include sensor/images for healthcare requiring low-latency rendering with high bandwidth for data stream upload [C-19].

Cloud software needs to be aware of the capabilities and performance of the underlying network infrastructure and be able to adapt with contingency plans to react to network outages and performance degradations [C-34] with the mechanisms to manage NaaS capabilities (SDN, NFV) to bypass problems in the telecommunications network, assuring and guaranteeing the operation of the cloud-based applications/services [C-26] [C-58]. Further work on adaptation is required on automated and secure load sharing in federated data centres with performance identification/monitoring functions considering the dependencies of cloud-based software on networking technologies [C-46]. Also vice-versa, the characteristics of the physical network and the communication fabric itself need to adapt automatically to the characteristics and performance metrics of the running cloud services [C-53]. Overall, vertical cooperation between networks, data centre infrastructures and application management is required to ensure end-to-end service continuity for users and business assurance for organisations [C-58].

The workshop discussions identified that further innovation is required in real-time and reliable cloud infrastructures – including data centres designed for high performance computing – to match advances in network technologies delivering ultra-low latencies: 5G promises to reduce latency by five times compared to 4G networks, for example.

Finally, service orchestration systems should investigate intelligent assembly and composition of cloud applications from software components and services driven by desired levels of assurance for selected service qualities [C-49].

3.8 Energy Efficiency in Cloud Computing

As mentioned in section 3.3 energy efficiencies in cloud computing can be obtained by exploiting hardware and firmware advances in VM, IO, memory management, CPU and GPU virtualisation, multi-core multi-machine scheduling and by adopting lightweight virtualisation, and natively distributed low level middleware [C-37] [C-58]. Further efficiencies can be made by incorporating environmental awareness into software development, where software takes into account its energy use, optimising computational complexities of cloud-based systems [C-46] based on generic metrics, benchmarks, and energy models [C-48].

Cloud computing can help realising greener computing practices by limiting power in end-devices as offloading processing to cloud resources means that a full-fledged workstation is not necessary anymore [C-2] [C-33]. Well-designed applications for the cloud should auto-scale and adapt to actual demands, using only the amount of resources that are necessary to comply with the service level objectives [C-2]. Advances in cloud federation will enable cloud service providers to exchange excess capacity with one another to meet demand and maximise profits and increase energy efficiency [C-51].

3.9 Managing Clouds and Services

More research is needed on self-managing and autonomic services running within the cloud with capabilities of self-healing and dynamic recovery [C-32] [C-46]. This includes the evolution of management APIs [C-58] and techniques for the efficient management of properties such as performance, availability and security within the cloud in light of the continued expansion of the cloud market both in Europe and globally [C-56].

Algorithms and procedures are required for evolving DevOps environments towards the automation of management patterns for deployment, cloning, offloading and reconfiguring systems for optimising cloud-based system organisation and processes [C-58].

3.10 Cloud Application Assembly and Automation

Current languages and development frameworks are not suitable to develop applications that can migrate to any device, nor exploit the parallelism of distributed, heterogeneous clouds [C-54]. Tools incorporating business support services to streamline and automate “concept-to-market” processes are required for assembled applications in the cloud enabling automated instantiation, deployment, configuration and life-cycle management multiple targeted cloud platforms (e.g. CloudStack, OpenStack, Amazon EC2) [C-49].

Algorithms and tools are required for the orchestration and interworking of autonomic cloud-based software resources/components, covering: dynamic context adaptation; advanced interworking using automation, adapter generation and service ontologies; model driven design, deployment and management for heterogeneous clouds; and large-scale scheduling [C-22] [C-49] [C-58].

Cloud computing facilitates run-time, dynamic adaptation of cloud-based applications in response to contextual changes and system failures. Analysis of dynamically available run-time data can assist in the identification of opportunities to adjust systems during live operation, for instance on decisions to scale cloud-based systems horizontally or vertically [S-10].

3.11 Cloud Support for Big Data Analytics and High Performance Computing

Tools and techniques are required for the design of cloud-based applications for big data analytics applications such as medical image processing and diagnosis with real-time performance and dealing with heterogeneous data sources [C-14] [C-19]. Models are required for a better understanding specific types and large quantities of data (e.g. TB/day) for optimising its compression and transmission over the network and between distributed components in the cloud. [C-44] [C-46]. To realise the full potential of cloud computing, particularly in the context of HPC and big data, a holistic approach to self-adaptation is needed to achieve the best possible performance while ensuring the most efficient use of resources, saving power and handling faults [C-53].

The post-consultation workshop identified that research into big data in the context of cloud computing should focus on aspects particular to cloud infrastructures, such as techniques and methods for partitioning data to optimise its processing in distributed clouds, to distinguish the identified priorities from more general big data analytical techniques funded by unit G3.

3.12 Cloud-based Vertical Markets and Applications

Several contributions identified the need to create building blocks to support cloud-based vertical markets/applications for interworking between sector actors [C-8]. Specific markets include: healthcare [C-19], finance, logistics, tourism [C-39] and education. The European education and research community is a well-suited vertical to act as an early adopter [C-50] deploying tools for managing digital content for online/offline learning with collaboration between educational institutions [C-35].

3.13 Cloud Experimentation

Platforms are required to experiment with cloud services and deployment configurations in realistic environments for both research and business objectives in configurations ranging from highly instrumented and controllable classical central data-centres, to highly innovative platforms of highly distributed clouds for edge computing. Flexible platforms will enable experimenting at scale with decentralised and federated approaches, investigating the performance of advanced hardware and firmware features. Experimental platforms will enable data collection for a range of metrics concerning the infrastructure: network, storage elements, hypervisor, characterising the QoE under a range of perturbations encountered on typical clouds under a wide range of load and fault conditions [C-47] [C-48] [C-50]. Integration of experimental platforms with international initiatives in, for example, the USA, Latin-America or Asia will allow experimentation at large scale [C-50].

It was reported during the post-consultation workshop that the Global Cloud Experimental Facilities workshop (8 October 2014, Florence, Italy) concluded that experimental cloud platforms should: attract applications that can be rapidly deployed on the market; investigate multi-cloud bursting with brokerage; cover federation of PaaS and SaaS; allow experimentation over inter-domain infrastructures; support services in different contexts; involve international testing, aligned with the IEEE InterCloud initiative; and support for testing and validation of services developed by SMEs.

There was unanimous support, during the post-consultation workshop, for testing cloud applications and services beyond the capabilities of current commercial cloud providers. Existing platforms, including those from GÉANT, Bonfire/FIRE, the FI-PPP and GRID 5000 in France, focus on different aspects of cloud computing experiments and with appropriate extensions they could support future experimentation requirements. Attention should be placed on ensuring portability between experimentation platforms. It was identified that further discussion is needed to identify the longer-term (five years into the future) features and technologies to be supported by platforms for experimentation, beyond the existing infrastructures in the EU and those available in the USA/Asia.

3.14 Simulations of Cloud Service Deployment

Tools are required, during early design and development phases and prior to actual deployment, to simulate the behaviour of cloud applications to assess performance under a variety of deployment configurations against non-functional requirements in the selected infrastructure, thereby aligning IT infrastructure resources with business objectives reducing development costs [C-8] [C-33] [C-48] [C-52]. In addition, simulations tools are needed to assist with dynamic reconfiguration strategies and scaling algorithms of distributed applications under various demand scenarios to take advantage of the unlimited resources provided by both single and federated cloud systems [C-33].

3.15 Standards for Cloud Computing

Open standardisation of service, platform and infrastructure interfaces is required to avoid the need for systems to comply with multiple technology-, vendor- and platform-specific APIs for data storage and access, identity management, access control, resource description and consumption models in both single and federated cloud systems [C-44] [C-47]. Several contributions identified the need for de-facto open standards rather than de-jure standards to speed up adoption [C-47] [C-54].

Clearly defined APIs enabling open development customisation, to publish cloud automation libraries, service models, best practices and policies are required to facilitate the uptake and utilisation of innovations by the cloud ecosystem as well as to advocate these by participating in open standards and interoperability activities including, for instance, OASIS CAMP and TOSCA, OCCI and CDMI [C-47] [C-49].

In addition to standard APIs and interfaces for interaction with cloud infrastructure, service/application specific standards need to be established for the widespread uptake of cloud services: such as interfaces for the semantic description and formal exchange of public records amongst public bodies [C-40].

The post-consultation workshop identified that because the cloud computing industry is largely driven by competition, the major cloud providers are unlikely to prioritise efforts on standardising APIs and interfaces. Standardisation of APIs for customer access and for federated cloud interworking is therefore a key area for attention.

3.16 Increased Uptake of Cloud Computing

In order to make cloud computing applicable and useable to the widest audience there must be tools and infrastructure that abstracts the user away from the computational complexities of the cloud [C-2] [C-18] [C-32]. There is a steep learning curve for engineers to understand the implications of deploying software via the cloud [C-22] and further automation is required in the integration, deployment and lifecycle management of software applications/products [C-18]. Training and exchange of expertise and knowledge transfer between institutions producing and consuming cloud services, including dissemination of good practice cases will increase awareness and expertise [C-50], especially if training is based around well-accepted and implemented standards [C-2].

Other factors to increase the adoption of cloud computing include the definition and adoption of standards and recommendations that will facilitate the migration of applications and services to cloud infrastructures [C-58], and the adoption of cloud computing by EU public administrations, potentially with the creation of a European G-Cloud [C-57]. Furthermore, future

H2020 projects will be more effective in increasing the uptake of cloud services by focussing on introducing new capabilities in cloud systems as cloud solutions are applied to new use cases and environments, rather than duplicating the work that has already been carried out by the private sector [C-56]. Further communication around the Cloud for Europe initiative should take place to ensure adoption and realisation of the benefits of the European Commission's investment on the issue [C-56].

In addition, increased trust in cloud computing by the measures identified in section 3.2 will increase the uptake of cloud by new sectors [C-6] [C-52].

Finally, new tools and programming models to assist with software migration from multi-core to the cloud [C-37] will assist with the conversion of legacy solutions and the uptake of cloud technologies [C-35]. Current languages and development frameworks are not suitable to develop applications that can migrate to any device, or exploit the parallelism of multiple computing resources [C-54].

3.17 Business Models

Research on new business model for cloud computing is required to aid migration and cloud uptake and to model the interactions between stakeholders in federated and distributed clouds. Issues include: the creation of ecosystems to migrate whole sectors to cloud infrastructures and services [C-58]; identification of market characteristics to increase the uptake of cloud computing [C-57]; access to resources on demand, whether directly in an inter-cloud model or mediated through in-house or third party brokers opens up possibilities for a multitude of novel business models, for example, cloud service providers will be able to exchange excess capacity with one another, to match supply and demand, increase profits and improve energy efficiency [C-51]; modelling value networks enabled by new actors such as a cloud software and services store to develop reference models from orchestration to revenue sharing for all stakeholders affected by the innovations [C-49].

Disruptive business process models and system architectures are needed to support the seamless reuse of public services from one smart city to another. At the process level, service compositions will require the ability to reason about system properties without an intimate knowledge of the implementation. At an architecture level, loosely coupled solutions will be sought allowing the agile and efficient composition of reusable public services with others that are produced in a local context [C-43].

SMEs are often involved in creating solutions to tackle specific business needs in the cloud computing ecosystem. However a better integration of data and processes produced by SMEs will avoid inefficiencies such as the storage of the same data on different platforms. Novel business models will enable more collaborative solutions and integration of business processes for more efficient operation, allowing individual SMEs to concentrate on their specialisation in the business chain rather than attempting to do everything or reinvent the wheel [C-12] [C-35] [C-51]. Adoption of new collaborative business models will aid European SMEs in gaining an EU-wide/global market reach for their technologies and products [C-18]. Involvement of SMEs will also be increased by the wide adoption of open standards for cloud deployment and portability to create a level playing field for small players that do not have the resources or market dominance to maintain an industry (or pseudo) standard and, thereby, enable them to compete in an open and standards-compliant marketplace [C-46] [C-47].

Furthermore, the workshop discussions identified that business and payment models for cloud services need further attention in the context of user expectation that many Internet- or cloud-based applications/services should be free and/or open source.

3.18 European Clouds

Several contributions identified a dependency on large cloud providers from outside of Europe. Establishing effective federations of public cloud providers within the EU would allow data (and revenue) to remain within the EU and be subject to European standards, certification and legislation [C-9] [C-15]. Furthermore this infrastructure could be tailored to the needs of European industry and scientific research with methods for splitting load between providers to optimise usage to facilitate high performance computing and big data analytics [C-52].

3.19 Regulation and EU policy

Several contributions referred to regulation and EU policy aspects of cloud computing, calling for a legal framework on security and data and related international treaties [C-15] [C-34] to prevent governments from accessing private and corporate data stored in cloud computing platforms without a specific warrant [C-17]. The workshop discussions concluded that an overall legal framework is required to regulate where, how and by whom data is stored, processed and accessed and suggested that Europe would be in a unique position for innovation if this was established together with the appropriate technical infrastructure.

An EU-based cloud infrastructure was proposed by the contributions mentioned in section 3.18, but, with appropriate international negotiations this could be extended to other regions to offer fast access around the world with similar safety and privacy guarantees and formal notification if, due to criminal or suspicious activities, personal/corporate data is to be shared with overseas agencies [C-25].

The workshop discussions concluded that common digital identity mechanisms would make authentication and access control in federated and distributed clouds spanning national borders more efficient for users and service providers alike and that policy makers should be encouraged to create conditions for the establishment of systems for EU-wide digital IDs.

Two contributions representing authors, artists and performers requested that the EU adopts appropriate legislation to establish private copying remuneration for cloud-based copyrighted media [C-29] [C-45].

Other aspects of EU policy/regulation raised by the consultation include: the requirement for cloud providers to provide traceability tools for data access [C-9]; cloud providers to prevent access to data by anyone other than the owner [C-16]; a strategy or policy addressing the rights of EU citizens opting for cloud services outside Europe [C-31]; financial measures to favour EU-based community, private or public clouds [C-31]; a legal framework to prevent overseas authorities without warrants from demanding access to EU citizens' data through their domestic legal frameworks without being granted via the proper legal channels [C-34]; prevention of research projects from data mining non-anonymised data [C-38]; legal changes should be considered in order to avoid blocking the advances in cloud and big data for certain interesting fields (affecting real life) such as personalised medicine or genomics, while guaranteeing the privacy of the information [C-44]; development of a European search engine based on anonymity and privacy which better matches European ideals, principles and laws than established popular search engines [S-4]. However the right balance needs to be struck between an over-regulated but safe cloud market and a vibrant cloud market [C-57].

SECTION 4. DETAILED ANALYSIS OF RESEARCH CHALLENGES AND PRIORITIES IN SOFTWARE (INCLUDING OPEN SOURCE)

The following sub-sections present the key research challenges and priorities for the H2020 ICT Work Programme 2016-17 on Software (including Open Source) as proposed in the written contributions and discussed in the workshop. The first four sections are broadly organised according to the software development lifecycle: section 4.1 covers software development methodologies and tools; section 4.2 extends this with specific issues of software for cloud computing; section 4.3 covers software testing methods and tools; and section 4.4 considers the challenges associated with software system adaptability. Section 4.5 addresses the challenges associated with open source software and section 4.6 covers software standards.

4.1 Software Development Methodologies, Tools and Techniques

Agile/scrum development methodologies can reduce development costs, increase customer satisfaction and adapt to volatile requirements changes. However, developing mission-critical and safety-critical systems with agile/scrum techniques requires further study as the existing tools and expertise for the development of mission-critical and safety-critical software are not suited for agile development methods [S-3].

Further attention should be put on requirements engineering rather than purely concentrating efforts on design and implementation issues [S-5] [C-59]. This is even more important in the case of multi-partner projects including dozens of stakeholders [S-17]. Software developers should be able to concentrate on defining the essential characteristics of systems which include domain models and data structures, algorithms for data processing and behaviour as seen by users and external systems. New techniques and tools are required for capturing requirements and compiling them to final systems. Details of underlying technological platforms should be hidden to allow engineers to concentrate more on early phases of the software development lifecycle [S-5].

Use of natural language for the collection and management of user stories for requirements capture is an easy-to-use tool for users but it is inherently ambiguous and too unstructured for developers to accurately map requirements to software systems. Research into new methods of requirements capture should be multidisciplinary and include Human Sciences experts; the adoption of Computational Linguistics tools could be a good starting point. Tools should be linked with CASE systems to enable a more formal input of captured requirements into the software engineering process [S-3].

Higher-level, intention-driven programming models for software replacement, recomposition and adaptability are required [S-13]. Goal oriented and data-intensive analysis techniques are needed to engineer the requirements for self-adaptive systems, to analyse user feedback and to model risks associated with the technical solutions [S-9].

Further work is required on software monitoring and measuring systems and metrics as current tools are based on complexity metrics developed for traditional software development processes and are not well suited to modern agile/scrum frameworks [S-3].

Tools and mechanisms for specifying, quantitatively assessing and strategically managing software quality are required to address time-sensitive business opportunities, fulfil market needs and avoid making losses on refactoring software to address quality deficits. Communication between stakeholders is important to maintain the link between externally perceived quality (e.g. software maintainability), internal characteristics of specific software artefacts (e.g. quality of software architecture), and software development activities (e.g. refactoring). Maintaining visibility and traceability of quality levels from different perspective will facilitate plans for short- and long-term quality improvement activities based on cost/benefit trade-offs [S-12].

End-users are increasingly providing feedback through reviews, ratings and comments in online forums, app stores and social networks. New research is required into software quality monitoring systems to make use of contextual data collected at runtime through various sensors and online data sources, smart objects/appliances. Methods and tools for gathering end-user feedback need to scale and enable automated analysis of feedback and contextual data to support requirements analysts, system architects, developers and project managers in decision-making tasks [S-15]. Furthermore, monitoring and analysis of user feedback combined with new tools to facilitate faster innovation cycles will enable personalisation of software to meet contextual user requirements and expectations, and enable forecasting of future user needs [S-15].

There is a wealth of data available on developer forums, forges, blogs and social networks. In addition, monitoring tools are available to extract information from cloud applications, services and web applications on the usage of program features, memory consumption of components, bandwidth for communication, etc. Big data analytics can be used to gain customer insight on user trends and preferences, to identify application performance and resource consumption metrics, to identify feature and performance improvement opportunities and the causes of software failures [S-10] [S-19]. Use of big data analytics should be further investigated for advancing methods and tools for software design, development leading towards automated software repair, component-based program synthesis and specification-driven optimisations [S-21].

4.2 Software Development for the Cloud

This subsection addresses the software engineering challenges for Software Development raised by the contributions to the consultation on Software (including Open Source). There is a related set of challenges in section 3.3 on Software Engineering for Cloud-based Systems as submitted to the consultation on Cloud Computing.

The ultimate goal for software engineering for the cloud is to create development environments that support the concept of “the cloud as the computer” meaning that an application or service running on a large number of distributed resources can be developed as easily as if it were running on a single-processor computer while being able to take advantage of scalable processing power and memory and being resilient to failures. Development environments and toolkits are needed to treat the cloud as a single virtual machine abstracting away from current development concerns such as how to distribute code and make it performant in distributed and heterogeneous environments [S-10].

As more systems are developed for cloud computing there is a shift from a centralistic, singular and static approach to software systems to a distributed, concurrent and adaptive one. Longer-term foundational and mid-term semi-foundational advances are required to identify new algorithmic and processing models for distributed, heterogeneous execution environments, for example processing without immediate access to data. Longer-term foundational advances require a coordinated effort, either in a single long-running project or through coordination between concurrent and/or sequential projects [S-13].

Software development tools are required for high quality, reliable cloud applications. Tools such as Erlang/OTP, developed for small scale, fault-tolerant distributed systems, need to be adapted for the development of larger-scale, distributed cloud applications, incorporating mechanisms to control the underlying network and processing resources according to the design patterns of cloud-based applications and services [S-24].

New approaches are needed for engineering socio-technical cloud-based software, such as systems for healthcare, emergency and disaster management, sustainability, smart cities, supporting organisational work, e-commerce and e-business collaboration and teamwork and social networks. Issues to be addressed include how to express accountability relationships formally: recognising accountability relationships as first-class software abstractions is the first and essential step toward systematically building socio-technical systems. [S-18]

Porting legacy software to the cloud raises many issues such as interactions with users and end devices, security, and the management of the transfer of large amounts of data between users and cloud-based processing/storage systems [S-7] [S-13]. More support is required in terms of tools and methodologies to facilitate migration and the uptake of cloud-based software.

4.3 Software Testing Methods and Tools

Model-driven software engineering tools are required to aid the automation of interoperability and compliance testing and support deeper reasoning to identify interoperability errors for cloud-based applications and services. Model-driven tools are also required to develop and manage cloud broker services that provide interoperability services between clients and vendors. Tools should cover both functional and non-functional mediation requirements, for example, preserving privacy and trust. Further work is required on the automatic generation of lightweight interoperability models based on deployed solutions, removing the need for developers to create models based on, for example, WSDL and BPEL. Techniques for using these models for generating software implementations need further research [S-6].

Automated testing tools are required to empower developers with the means of conducting extensive, unattended validation, targeting functional and non-functional (e.g., security) properties of the software and focusing on the quality of the test scenarios [S-9].

4.4 Dynamic Adaptation of Software Systems

Conventional information system design and analysis approaches are not well suited to managing risks in evolving systems based on dynamic agile system engineering approaches for software definition, composition and deployment. Approaches based on expert analysis cannot be scaled and sped up sufficiently to be used in large-scale, autonomically provisioned infrastructure and applications. New approaches are needed based on automated risk identification and analysis, detection and mitigation that can be applied, at scale, during design and operational phases of distributed applications and services [S-11]. The European software industry needs to develop new value-driven and adaptive real-time business paradigms and new technical infrastructures and tools are required to support this transformation [S-19].

Despite the proliferation of semantic web services there is currently a gap for new intelligent software clients to autonomously consume and compose web services given a specific goal [S-22].

4.5 Open Source

Collaborative research and innovation projects typically deliver proofs of concept, demonstrations and components rather than market-ready offerings. If these are released as open source systems then they can be transformed into technology commons for the creation of value by third-party users, systems integrators and software/service vendors. The objective would be to significantly increase adoption of collaborative projects' open source software outcomes. This would involve enforcing open source governance mechanisms and procedures, developing a dedicated marketplace for the sourcing of open source technology commons, creating business opportunities through government IT spending power and leveraging global EU-centric communities [S-20]. The post-consultation workshop discussions concluded that an EU-based open source community would be a useful mechanism to assist with the exploitation of the results of EU-funded open source projects, and that software publication and engagement with open source communities should be undertaken throughout their lifetime of the project and not just at the end.

There is a need to manage risks in the adoption and integration of open source cloud platforms, including licensing issues, mismatched expectations and code vulnerabilities [S-26]. Further research is required on the provision of quality testing platforms for open source software configuration management scripts and auto-deployment bundles (possibly using container technology) and the establishment of an independent marketplace or trusted forge for quality-assured SCM scripts and/or auto-deployment bundles to overcome quality issues of today’s scripts and bundles. These quality issues include operating system heterogeneity, hidden dependencies, installation mechanisms, SCM tool version, and underlying platform incompatibility. The marketplace should be independent and semi-curated including reputation metrics based on ratings and other user feedback [S-14].

Adopting Open Web standards is especially important for open source related projects to encourage reuse [S-2].

For certain application domains such as Smart Cities the reuse of open source software combined with novel collaborative methods, techniques, tools and platforms for producing public services may help to dramatically increase the level of software reuse, while decreasing public service software production costs [C-43]. Work on specific European open source applications, such as office applications could save the EU significant money on proprietary software licences [S-23]. Migration of public sector systems away from proprietary software to open source would be a key driver for the wider adoption of open source software [S-8].

The workshop discussions identified that although almost all popular cloud-based systems (Facebook, for example) are heavily based on open source software, there are ongoing trust concerns with deploying open source systems and services for enterprises. A better understanding of the concerns is required, together with technical solutions for addressing the trust issues, in order to increase the uptake of open source software by more industrial and commercial sectors.

4.6 Standards for Software

Innovation and standardisation, in the ICT software area, can take place almost simultaneously as the timeframe from the first definition of an API to access a new device, for example, to its inclusion in a global standard such as HTML5 or CSS3 has been reduced to months/years compared to a few decades in the past standardisation world. This makes it possible for research projects to have a greater impact on standards within their lifetime. The principles of open standardisation, such as those defined by the Open-Stand.org initiative, should be taken into account by future research projects; either for the standards they use or for those they plan on extending or improving. Support actions in the area of ICT standards would promote the wider adoption of open standards, encourage participation in the standardisation process and organise training for researchers on the use of open standards [S-2].

SECTION 5. BACKGROUND AND TERMS OF REFERENCE

The *Software & Services, Cloud* unit (DG CONNECT E.2) launched an open web-based public consultation on 10 September 2014 on the research and innovation challenges in Cloud Computing and Software (including Open Source). A set of reports were identified by the EC to trigger discussions (see section 7) and all interested sector actors – from industry, academia, research institutes, SMEs and user communities – were invited to comment on the topics raised by the reports and to identify further key research themes and priorities for the Horizon 2020 Work Programme 2016-17. A post-consultation workshop was held on 4 November 2014 where the written contributions were presented and discussed to identify research topics for the upcoming Work Programme and to ensure the participation of all interested stakeholders.

Sixty-one written contributions on Cloud Computing and twenty-six on Software (including Open Source) were submitted during the consultation period (see section 5). Twenty-seven presentations were made at the workshop: sixteen on Cloud Computing and eleven on Software (including Open Source). The workshop was organised around five sessions as follows (the full agenda is contained in appendix A):

- Federated and Distributed Clouds, Interoperability (five presentations)
- Other Research Challenges in Cloud Infrastructures and Services, part I (six presentations)
- Other Research Challenges in Cloud Infrastructures and Services, part II (five presentations)
- Software Methods and Tools (eight presentations)
- Open Source Software, Interoperability and Standards (three presentations)

The following charts show a breakdown of the written submissions by country and by stakeholder sector for contributions on Cloud Computing.

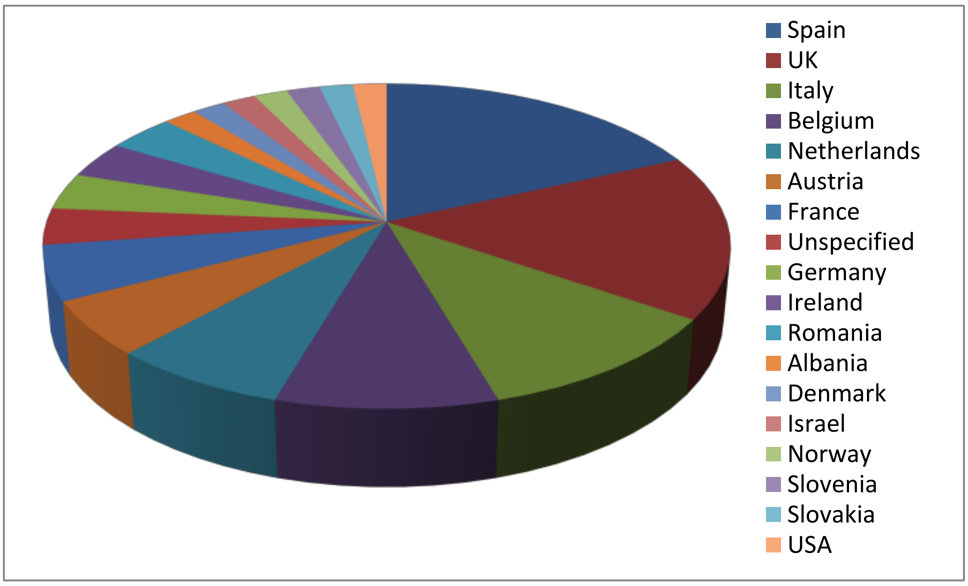

Figure 1. Contributions on Cloud Computing by country

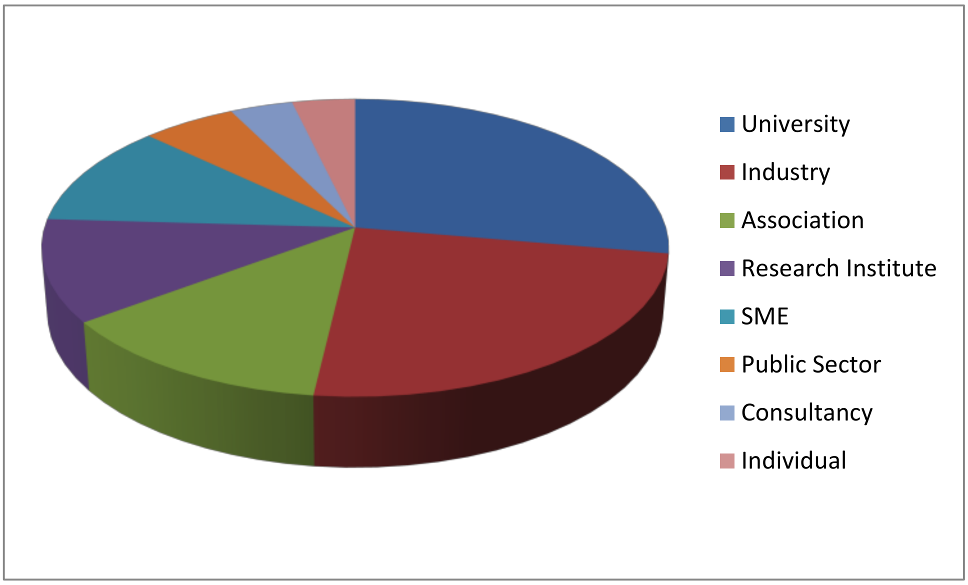

Figure 2. Contributions on Cloud Computing by organisation type

The following charts show a breakdown of the written submissions by country and by stakeholder sector for contributions on Software (including Open Source).

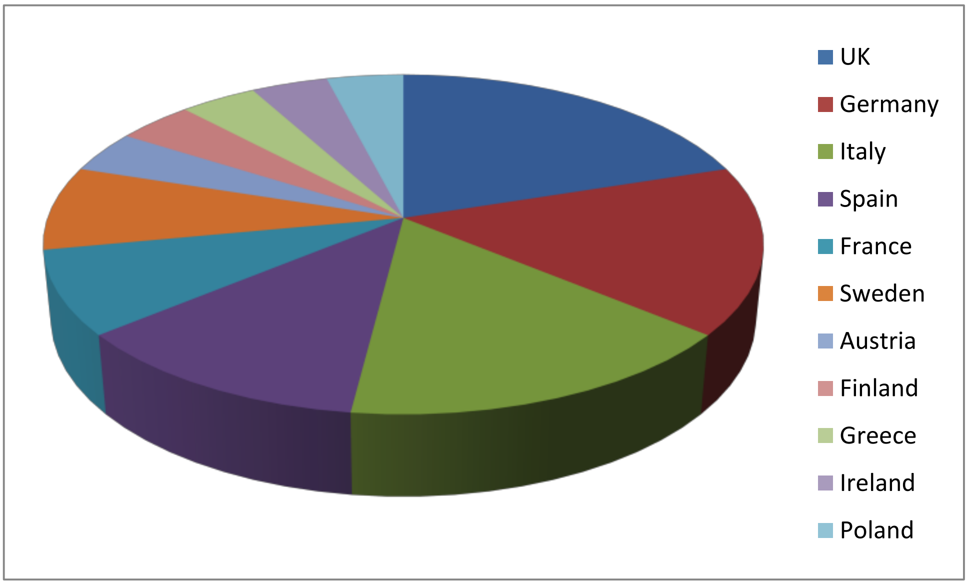

Figure 3. Contributions on Software (including Open Source) by country

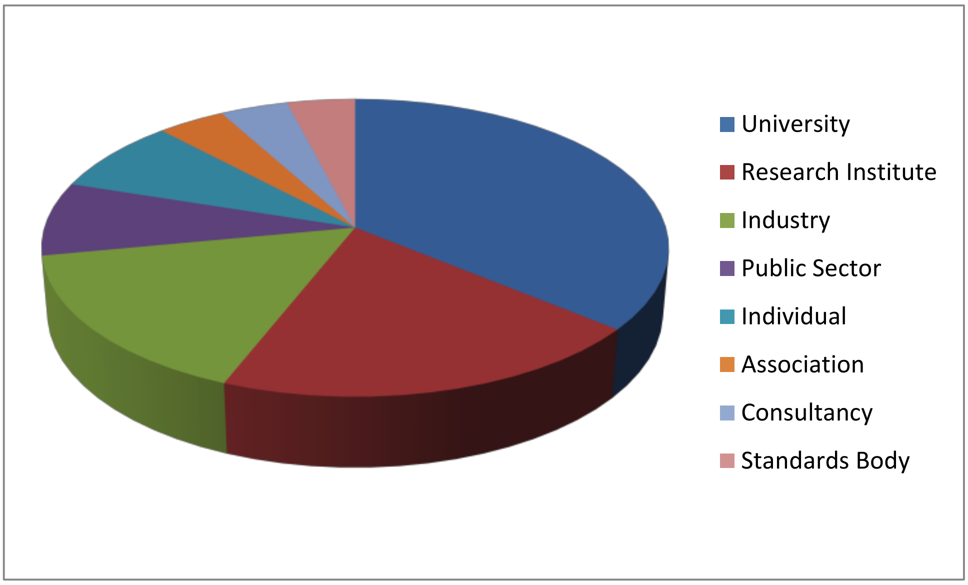

Figure 4. Contributions Software (including Open Source) by organisation type

As shown by the charts, the spread of contributions to the consultation process between countries and organisation sector was broadly in line with the size of EU member states and with the involvement of sector actors in current programmes.

SECTION 6. WRITTEN CONTRIBUTIONS TO THE CONSULTATION PROCESS

Copies of the written contributions are available at:
https://ec.europa.eu/digital-agenda/events/cf/cloud-computing-software-engineering/stream-items.cfm?id=277&by=276

Cloud Computing

[C-1] *Contribution to Consultation on Cloud Computing*, Francisco MINGORANCE, Europa Insights sprl, Belgium.

[C-2] *Reflections on Cloud computing research*, Jose MIGUEL-ALONSO, University of the Basque Country UPV/EHU, Computer Architecture and Technology, Spain.

[C-3] *ITC*, Lirim BYLYGBASHI, University of Prishtina, IT, Albania.

[C-4] (duplicate contribution)

[C-5] *Cloud Computing in the age of DevOps*, Eliot SALANT, IBM Haifa Research Lab, Israel.

[C-6] *Cloud vs on-premise*, Daniel ROS, Terrassa City Council, Spain.

[C-7] *Semantics in and for Cloud Computing*, Andreas HARTH, Karlsruhe Institute of Technology, AIFB, Germany.

[C-8] *Cloud Public Consultation - My views*, Agustin GONZÁLEZ-QUEL, Ariadna Servicios Informáticos, Spain.

[C-9] *To support Public Cloud European Users*, Cyril BARTOLO, EUROCIO, France.

[C-10] *Cloud Brokerage of cloud services*, Rastislav NECZLI, ITMG, s.r.o., Slovakia.

[C-11] *Smart Sensing in Heterogeneous Clouds for Smart Living*, Antonio PULIAFITO, University of Messina, Computer Department (DICIEAMA), Italy.

[C-12] *SME's need to integrate*, Maurice VAN DER WOUDE, BPdelivery B.V., Executive board, Netherlands.

[C-13] *PaaS Software*, Mario MEIR-HUBER, IDC, Austria.

[C-14] *Cloud and Big Data for European industries*, Mario MEIR-HUBER, IDC, Austria.

[C-15] *Concerns for privacy*, Benjamin MARCH.

[C-16] *Privacy and Anonymity*, P MAUD, LGSS, UK.

[C-17] *Some minimal standards for cloud computing*, Jeremy WICKINS.

[C-18] *Increase cloud automation for European businesses and Governments*, Paola PONTICELLI, Liberologico, Bid and R&D, Italy.

[C-19] *Clouds from a healthcare device perspective*, Frank VAN DER LINDEN, Philips Healthcare, CSIO office, Netherlands.

[C-20] *Dynamic system engineering*, Mike SURRIDGE, IT Innovation Centre, UK.

[C-21] *Data and Information Quality in the Cloud*, Adam TRENDOWICZ, Fraunhofer Institute for Experimental Software Engineering, Measurement, Prediction, and Empiricism, Germany.

[C-22] *Concerns for how software is distributed through the Cloud*, Mariam KIRAN, University of Bradford, Computer Science, UK.

[C-23] *Creating a better cloud environment in Europe and building trust*, Peter WITSENBURG, Belgium Cloud, Belgium.

[C-24] *Towards a Grand Coalition for Trustworthy Cloud Computing*, Theo LYNN, Dublin City University, Irish Centre for Cloud Computing and Commerce, Ireland.

[C-25] *EU Cloud Privacy and Integrity*, Tahir KHAN, Agilent, USA.

[C-26] *Contribution of the CEPIS LSI - Statement on Cloud Computing Security and Privacy Issues*, Marko HÖLBL, University of Maribor, Faculty of Electrical Engineering and Computer Science, Slovenia.

[C-27] *Accountability-Based Software Engineering*, Amit CHOPRA, Lancaster University, School of Computing and Communication, UK.

[C-28] *Cloud Computing: Need for a new generation of standards*, Luca NICOLETTI, SOGEI S.p.A., DVE-AST-SAS, Italy.

[C-29] *Joint statement of creators and performers organisations*, Nicole SCHULZE, AEPO-ARTIS, Belgium.

[C-30] *Weak data consistency in cloud storage systems*, Alexey (Olexiy) GOTSMAN, IMDEA Software Institute, Spain. (Co-authored by Jade ALGLAVE, Microsoft Research and University College London, UK).

[C-31] *Critical concerns regarding the future of cloud computing in EU*, Bogdan Alexandru CHELARIU, North-East Regional Development Agency Romania - Brussels Office, Belgium.

[C-32] *Considerations for Cloud Computing*, Stephen MCGOUGH, Durham University, School of Engineering and Computing Sciences, UK.

[C-33] *Challenges for multi-cloud applications deployed on heterogeneous environments and new cloud business models*, Leire ORUE-ECHEVARRIA, TECNALIA, ICT / European Software Institute Division, Spain.

[C-34] *Security*, Torben RYBNER, Ziberex, Denmark.

[C-35] *Strengthening the industrial structure of SMEs*, Eduardo SÁNCHEZ, AMETIC - ASTI, Spain.

[C-36] *Improving competitiveness in education processes*, Eduardo SÁNCHEZ, AMETIC - ASTI, Spain.

[C-37] *Efficient use of cloud infrastructures*, Tomás F. PENA, University of Santiago de Compostela, CITIUS - IT Research Centre, Spain.

[C-38] *Privacy by Design for cloud, new research directions*, Alboaie SÎNICĂ, Axiologic, Software, Romania.

[C-39] *Improving the competitiveness of the tourism sector*, Eduardo SÁNCHEZ, AMETIC - ASTI, Spain.

[C-40] *eDocument interoperability: Semantic transformation of administrative files*, Helmut FALLMANN, Fabasoft, Austria.

[C-41] *Secure Personal Data and Documents Vault: Secure electronic safe*, Helmut FALLMANN, Fabasoft, Austria.

[C-42] *High quality reliable cloud-native applications*, Simon THOMPSON, University of Kent, Computing Laboratory, UK.

[C-43] *Cloud Public Services for Smart Cities: the way ahead*, Xavier FRANCH, Technical University of Catalonia (BarcelonaTech), Service and Information System Engineering, Spain.

[C-44] *Cloud and scientific data management and computation*, Alfons NONELL-CANALS, Mind the Byte, S.L., Spain.

[C-45] *GESAC's contribution to the consultation on cloud computing*, Burak OZGEN, GESAC, Belgium.

[C-46] *TSSG Position Paper on Cloud Computing and Software*, Brian FOLEY, Waterford Institute of Technology, TSSG, Ireland.

[C-47] *EGI recommendations for the Horizon 2020 Work Programme 2016-2017*, Michel DRESCHER, Stichting European Grid Infrastructure, Netherlands.

[C-48] *Position paper from Inria*, Frédéric DESPREZ, INRIA, Research Department, France.

[C-49] *Cloud Evolution: Delivery through Horizontal Services*, Theo DIMITRAKOS, BT Innovate, Centre for Information and Security Systems Research, BT Innovate, UK.

[C-50] *GÉANT response to EC Consultation on Cloud Computing*, Andres STEIJAERT, SURFnet, Netherlands.

[C-51] *Cloud Computing towards the future: 451 Research*, Philip INGLESANT, 451 Research, UK.

[C-52] *Cloud Perspectives for next calls: more enlarging the scope*, Antonio CORRADI, Università di Bologna, DISI, Italy.

[C-53] *Supporting the Cloud Computing priorities report from the CloudWatch Concertation Meeting*, Ernst Gunnar GRAN, Simula Research Laboratory, Norway.

[C-54] *CLOUD Data Platforms and Programming Paradigm for Next Generation APPs*, Andrea MANIERI, Engineering Ingegneria Informatica S.p.A., R&D labs, Italy.

[C-55] *European Digital ID, European Cloud Based Key Storage, European Electronic Delivery*, Helmut FALLMANN, Fabasoft, Austria.

[C-56] *Hewlett-Packard (HP) Contribution*, Irena BEDNARICH, Hewlett-Packard, Government Relations, Belgium.

[C-57] *Response to the public consultation on future research and innovation priorities in cloud computing*, Nicky STEWART, Skyscape Cloud Services Ltd, Commercial, UK.

[C-58] *Nessi views on cloud computing challenges*, Valère ROBIN, Orange Labs, Products and Services, France.

[C-59] *Requirements engineering needs more tools!!!!,* José FUENTES, The REUSE Company, Commercial Department, Spain.

[C-60] *Improve predictability in Cloud-based scenarios*, Mauro MARINONI, Scuola Superiore Sant'Anna, Institute of Communication, Information and Perception Technologies (TeCIP), Italy.

[C-61] *Secure and High Assurance Cloud Computing*, Markus TAUBER, AIT Austrian Institute of Technology GmbH, Austria.

Software (including Open Source)

[S-1] *PAAS mature enough?,* Daniel ROS, Terrassa City Council, Spain.

[S-2] *W3C Input*, Daniel DARDAILLER, World Wide Web Consortium - W3C, Liaisons, France.

[S-3] *Adopting "Agile" methodology in mission critical software production*, Angelo MESSINA, Italian Ministry of Defence Army General Staff, Logistic, Italy.

[S-4] *European Search engine*, Alan LUKASZEWICZ, UK.

[S-5] *From Requirements to Cloud in a Snap?,* Michal SMIALEK, Warsaw University of Technology, Poland.

[S-6] *Model-driven interoperability*, Paul GRACE, IT Innovation, University of Southampton, UK.

[S-7] *Reflections on Cloud computing research*, Marco NATALI, IDS Ingegneria dei Sistemi, Italy.

[S-8] *Open Source as European Core in society*, Magnus LINDBERG, Sweden.

[S-9] *Contribution from Fondazione Bruno Kessler, Trento, Italy; Software Engineering research unit*, Paolo TONELLA, Fondazione Bruno Kessler, CIT, Italy.

[S-10] *NESSI's Complementary Recommendations for WP 2016/2017 on SOFTWARE ENGINEERING*, Andreas METZGER, University of Duisburg-Essen, Paluno (The Ruhr Institute for Software Technology), Germany.

[S-11] *Dynamic system engineering*, Mike SURRIDGE, IT Innovation Centre, UK.

[S-12] *Strategic Management of Software Quality*, Adam TRENDOWICZ, Fraunhofer Institute for Experimental Software Engineering, Measurement, Prediction, and Empiricism, Germany.

[S-13] *Challenges in Software Engineering*, H2020, Lutz SCHUBERT, University of Ulm, Germany.

[S-14] *Towards a Trusted Marketplace for Quality-Assured Open Source SCM Scripts*, Theo LYNN, Dublin City University, Irish Centre for Cloud Computing and Commerce, Ireland.

[S-15] *Software evolution driven by User Feedback and Contextual Data*, Anna PERINI, Fondazione Bruno Kessler (FBK), Trento, Italy, Center for Information and Communication Technology research center (FBK-ICT), Italy.

[S-16] *Coping with weak data consistency in the development of multicore and cloud software*, Alexey (Olexiy) GOTSMAN, IMDEA Software Institute, Spain. (Co-authored by Jade ALGLAVE, Microsoft Research and University College London, UK).

[S-17] *No Software Engineering without Requirements*, Christian WUENCH, HOOD GmbH, Germany.

[S-18] *Accountability-Based Software Engineering*, Amit CHOPRA, Lancaster University, School of Computing and Communication, UK.

[S-19] *Need for Speed*, Janne JÄRVINEN, F-Secure, Finland.

[S-20] *A Proactive Open Source Strategy for EU-Funded Collaborative Software Projects*, Cedric THOMAS, OW2, France.

[S-21] *Big Code*, Kyriakos CHATZIDIMITRIOU, Aristotle University of Thessaloniki, Electrical and Computer Engineering, Greece.

[S-22] *Semantic web of services*, Kyriakos CHATZIDIMITRIOU, Aristotle University of Thessaloniki, Electrical and Computer Engineering, Greece.

[S-23] *European Open Source Office Software*, Helmut FALLMANN, Fabasoft, Austria.

[S-24] *High quality reliable cloud-native applications*, Simon THOMPSON, University of Kent, Computing Laboratory, UK.

[S-25] *Transparent and Traceable Financial Information*, Lars G BERGLOF, TS Solutions AB, Senior management, Sweden.

[S-26] *Managing risks in the adoption and integration of open source cloud platforms*, Xavier FRANCH, Technical University of Catalonia (BarcelonaTech), Service and Information System Engineering, Spain.

SECTION 7. REFERENCES

The following reports were provided as input to the consultation process to trigger discussions. All documents are available for download at: https://ec.europa.eu/digital-agenda/en/news/public-consultation-cloud-computing-and-software

Cloud Computing

[CW14] *Shaping the H2020 LEIT ICT WP2016-2017: Recommendations from the CloudWATCH Concertation meeting*, 11 September 2014, DG Connect, Brussels, Belgium.

[EC12] *Unleashing the Potential of Cloud Computing in Europe*, Communication from the Commission to the European Parliament, the Council, the European Economic and Social Committee and the Committee of the Regions, Brussels, 2012.

[EC13] *European Council conclusions*, October 2013.

[EC14] *Towards a thriving data-driven economy*, Communication from the Commission to the European Parliament, the Council, the European Economic and Social Committee and the Committee of the Regions, Brussels, 2014.

[JEFFREY] *COMPLETE COMPUTING: TOWARD INFORMATION, INCENTIVE AND INTENTION, Research Priorities in Cloud Computing, in the context of Software and Services, taking into account Internet of Things, Future Internet and Big Data,* Keith Jeffery, Lutz Schubert.

[MCGEER] *The Distributed Cloud as the Future of the Cloud and the Internet*, Rick McGeer, Chief Scientist, US Ignite.

[NESSI-C] Cloud Computing, NESSI Recommendations for WP 2016/2017, NESSI, Networked European Software and Services Initiative, September 2014.

[ZYSMAN] *The Implications of Cloud for Integrated Research and Innovation Strategy*, John Zysman, University of California Berkeley.

Software (including Open Source)

[CW14] *Shaping the H2020 LEIT ICT WP2016-2017: Recommendations from the CloudWATCH Concertation meeting*, 11 September 2014, DG Connect, Brussels, Belgium.

[ISM] *Italian Software Engineering Manifesto*, Italian Association for Software Engineering.

[ISTAG] *Toward a Strategic Agenda for Software Technologies in Europe*, ISTAG Report on 2020 Key Software Technologies.

[NESSI-S] *SOFTWARE ENGINEERING, Key Enabler for Innovation*, NESSI White Paper, Networked European Software and Services Initiative, July 2014.

SECTION 8. ACRONYMS AND ABBREVIATIONS

API	Application Programming Interface
ASIC	Application-specific Integrated Circuit
BPEL	Business Process Execution Language
CAMP	Cloud Application Management for Platforms
CASE	Computer-aided Software Engineering
CDMI	Cloud Data Management Interface
CPU	Central Processing Unit
CSS	Cascading Style Sheets
DevOps	Development and Operations
DG CONNECT	European Commission Directorate General for Communications Networks, Content and Technology
EC	European Commission
EC2	Elastic Cloud Compute
ETP	European Technology Platform
EU	European Union
FI-PPP	Future Internet Public Private Partnership
FIRE	Future Internet Research and Experimentation
FP7	EU Seventh Framework Programme for Research
G-Cloud	Government Cloud
GÉANT	Pan-European research and education network interconnecting National Research and Education Networks
GPGPU	General Purpose Computing on Graphics Processing Units
GPU	Graphics Processing Unit
H2020	Horizon 2020
HPC	High Performance Computing
HTML	HyperText Markup Language
IaaS	Infrastructure as a Service
ICT	Information and Communications Technology
ID	Identifier or Identity
IEEE	Institute of Electrical and Electronics Engineers
IO	Input/Output
IoT	Internet of Things
IP	Internet Protocol
ISP	Internet Service Provider

ISTAG	Information Society Technologies Advisory Group
IT	Information Technology
NaaS	Network as a Service
NESSI	Networked European Software and Services Initiative
OASIS	Organisation for the Advancement of Structured Information Standards
OCCI	Open Cloud Computing Interface
OTP	Open Telecommunications Platform
PaaS	Platform as a Service
QoE	Quality of Experience
QoS	Quality of Service
R&D	Research and Development
RAMCloud	Random Access Memory Cloud
SaaS	Software as a Service
SCM	Software Configuration Management
SDN	Software Defined Networking
SLA	Service Level Agreement
SME	Small and Medium Sized Enterprise
TOSCA	Topology and Orchestration Specification for Cloud Applications
UI	User Interface
VM	Virtual Machine
WP	Work Programme
WSDL	Web Services Description Language
XaaS	Everything as a Service

APPENDIX A. WORKSHOP AGENDA

Copies of the presentation slides are available at:
https://ec.europa.eu/digital-agenda/en/news/public-consultation-cloud-computing-and-software

Time	Session Title
08:30	**Registration**
09:00	**OPENING SESSION: Consultation process for Cloud Computing and Software in the context of the LEIT-ICT work programme 2016-17 of Horizon 2020** • Francisco Medeiros, Acting Head of Unit, DG Connect E2
09:20	**SESSION 1: Federated and Distributed clouds, Interoperability** • Philip Inglesant, 451 Research • Maurice van der Woude, BPdelivery B.V. • Helmut Fallmann, Fabasoft • Leire Orue-Echevarria, TECNALIA • Luca Nicoletti, SOGEI S.p.A.
10:10	• Q&As and open discussion on session 1 topics
10:30	Coffee Break
10:50	**SESSION 2: Other Research Challenges in Cloud Infrastructures and Services (Part I)** • Valère Robin, Orange Labs • Alan Davy, Waterford Institute of Technology • Andres Steijaert, SURFnet • Andrea Manieri, Engineering Ingegneria Informatica S.p.A. • Frédéric Desprez, INRIA • Nicky Stewart, Skyscape Cloud Services Ltd
11:50	• Q&As and open discussion on session 2 topics
12:20	Lunch Break
13:20	**SESSION 3: Other Research Challenges in Cloud Infrastructures and Services (Part II)** • Yannick Legré, EGI.eu • Mariam Kiran, University of Bradford • Alfons Nonell-Canals, Mind the Byte, S.L. • Frank van der Linden, Philips Healthcare • Dana Petcu, West University of Timisoara
14:00	• Q&As and open discussion on session 3 topics

14:20	**SESSION 4: Software Methods and Tools** • Andreas Metzger, University of Duisburg-Essen • Angelo Messina, Italian Ministry of Defence Army General Staff • Marco Natali, IDS Ingegneria dei Sistemi • Paul Grace, IT Innovation – University of Southampton • Michal Smialek, Warsaw University of Technology • Alexey Gotsman, IMDEA Software Institute • Simon Thompson, University of Kent • Anna Perini, Fondazione Bruno Kessler
15:40	• Q&As and open discussion on session 4 topics
16:10	Coffee Break
16:30	**SESSION 5: Open Source Software, Interoperability and Standards** • Theo Lynn, Dublin City University • Cedric Thomas, OW2 • Xavier Franch, Technical University of Catalonia
17:00	• Q&As and open discussion on session 5 topics
17:20	**CLOSING SESSION: Next Steps and Closing** • Jorge Gasós, Head of Sector, DG Connect E2
17:30	End of the workshop