Recommendations for the Workprogramme H2020-ICT-2018-2019

Topics: Cloud middleware, Datacentre services, Cloud model evolution and Cloud service market orientation

Editor:

Dana Petcu, West University of Timisoara, Romania

December 24, 2015

Introduction

Aim

This open-access document aims to provide a vision about the future in research and development in the the period 2018-2020 in the fields related to Cloud models, middleware and market, as well as to datacentre services. Following this vision four recommendations are presented for the workprogramme H2020-ICT for 2018-2019.

Document ownership

The document is reflecting the opinions of several representatives of the projects involved in the cluster "New Approaches for Infrastructure services". The cluster intends to be a forum for discussing the current research and innovation challenges encountered at infrastructure-as-a-service level generated by the desire to improve the user experiences and the efficient use of the available resources. The current trends are including the integration of special devices from high performance computing ones to mobile devices, the design of decentralised service-oriented systems, the improvement of the virtualization technologies, the overcome of portability and interoperability issues, or the automation the organisation and management of the back-end resources. Cloud-based applications from the fields of Internet-of-Things and Big Data are expected to challenge the new services.

The projects that are member of the cluster are: AppHub, ARCADIA, CloudLightning, CloudSpaces, CloudWave, ClouT, BEACON, DICE, ENTICE, iKaaS, INPUT, IOStack, MIKELANGELO, Moble Cloud Networking, Mo-bizz, MODAClouds, MUSA, RAPID, SPECS, SWITCH. Details about the cluster activities, aims and results are available at:
https://eucloudclusters.wordpress.com/new-approaches-for-infrastructure-services/

Organization

The document is organized as follows. Four recommendations are provided for the following topics:

- Cloud middleware
- Datacentre services
- Cloud model evolution
- Cloud service market orientation

The visions related to the evolution in this topics are revealed after the recommendations.

Scope

The document do not include any recommendations related to the topics of Software Engineering, Inter-Cloud, Data Protection, Security and Privacy that are subjects of the other clusters.

Sources

In order to prepare this document a questionnaire related to more than 25 visions of the future in the above described topics reported in the literature has circulated by the cluster member in order to identify the ones that are expected to be of interest for 2018-2010. The current document reflects the selected visions following the majority of the 42 feedbacks received from the questionnaire. The questionnaire and the answers are available on the cluster site.

Recommendation 1 ***(related to Cloud middleware)***

Support the ***evolution of the Cloud towards its omnipresence****, freeing of the Cloud service consuming software from the Cloud services, by ensuring that the new generation software, or existing modular or event-reactive ones, will be able to be described to an abstract level that is service agnostic, will be able to form automatic and transparent combinations of hardware and software resources according to its needs, while resource provisioning, deployments, runtime migrations, multi-tenancy with cost-effectiveness and data protection, or the recovery from minute-to- minute failures will be managed automatically.*

This recommendation is based on the following ***visions****:*

The software is expected to float away from the hardware and software resources: the software will described in abstracted space and it goes through several adaptors before it interacts with hardware (being therefore practically hardware agnostic).

This trend is reflected in the recent Lambda architecture, in which code is executed in response to an event, ensuring an important step in moving away from server-centric design. The Cloud becomes a generic compute engine, and the developers do not need to organize the resources as they simple just run the code.

Applications will no more belong to one cloud, but will be able to reside in various clouds and to use various on-premise other applications. Different parts of applications will be able to float around in and out of current cloud, e.g. following a policy-based service-level agreement management plan. The software development process is expected to place an emphasis on modular software, allowing their components to be modified without shutting down the full.

The application should be able to form automatic and fleeting associations of hardware and software resources according to their needs, e.g. following a social-media style in which a database can ‘like’ a server or ‘like’ a storage array (this approach requires that the resources of a datacentre will mould itself around the task required, rather than following the current opposite way). Application developers will no longer need to worry about provisioning servers, storage, or communication as the provisioning process will happen automatically. New software systems need to be developed to deal with the likely minute-to-minute failures of the consumed resources.

One challenge ahead is the production of multi-tenant software by a software developer. When multiple users access the software relying on one Cloud, the software developer needs to make sure that data is kept separate and the associated costs are handled effectively.

Recommendation 2 ***(related to datacentre services)*****:**

Support the ***evolution of the software-defined datacentres as ecosystems****, in which services are abstracted from infrastructure, changes and updating are done automatically based on intelligent orchestration or new database tools, security is software-defined; such ecosystem should enable warehouse-scale computing using purpose-designed chips, new services like supercomputing on demand or massively federated, scalable software architecture with orchestration through network awareness.*

This recommendation is based on the following ***visions****:*

The software-defined datacentres with all infrastructure virtualised and delivered as a service is changing the way organisations view the value of the underlying physical infrastructure. Within the datacentre the services will be abstracted from the infrastructure.

Technologies currently limited to supercomputing will make it into the mainstream. The web-server-sized instances are replaced by on-demand private clusters. New cloud services are expected to emerge, like supercomputers on demand, high-performance storage, or new ways of storing and processing data.

Low-power processors will be able to treat many workloads in the highly automated datacentres to support massively federated, scalable software architecture. A new generation of warehouse- scale computing is coming and custom chips for cloud partners and purpose-designed chips will be soon a reality of the datacentre.

The datacentres will be like biological organisms: having different states, growing and shrinking according to workloads, automatically corrected and changed. The hardware will be controlled from a single point and an over-arching system will rule equipment via software.

When large-scale applications with thousands components are deployed across multiple datacenters the SLAs will be enforced and software usage and faults needs to be monitored and managed.

Cloud management and orchestration needs to be aligned with various products and services. Consequently, everything from orchestration to database tools will evolve. Datacentre operators can add value to cloud orchestration through network awareness and integration of cloud orchestration with their network management platforms. Software-defined security will become part of the software-defined datacentre.

Recommendation 3 ***(related to the Cloud model evolution):***

Support the ***evolution of the Cloud computing model towards the integration with machine-to-machine computing****, solving the challenges of separate technology stacks and dealing with limited memory, storage and computation capacity of the edge devices, speed of deployment, resource distribution, cost-effective scalability, resilience, easiness of management or security.*

This recommendation is based on the following ***visions****:*

Inverse cloud models, like machine-to-machine computing or geo-distributed cloud, have been recently proposed as an alternative to the bandwidth-intensive Cloud approach. Company branded versions, like fog computing, edge computing, in-network distributed cloud, or Cloud 2.0 were proposed in this context.

Machine-to-machine computing provides compute, storage, data or application services to client endpoints, like the Cloud computing model. Unlike it, machine-to-machine computing distributes data transmission among endpoints and routers, instead sending them to a server. However, it has its challenges, including limited resources, limited network capacity, security, or resource distribution.

The Fog computing promises to bring small computation and storage capabilities, supporting execution of applications that require low latency interactions with sensors, actuators or end-users. Fog nodes are implemented using embedded systems, in industrial control boards or home routers. Their limited memory, storage and computation is main challenge in their integration into the Cloud architecture to enable the execution of application logic. Their integration will leverage containers as virtualization technology for application delivery and execution. The success of Fog computing depends on the resilience of the smart gateways directing tasks on an Internet teeming with IoT devices. This smartening will rely on features such as out-of-band access, automatic detection and recovery from outages, cellular connectivity, or high-level security. Moving the processing of data to the edge raises the challenge to maintaining the availability of these smart gateways and their communication path to the cloud. Resilience is needed for business continuity, with redundancy, security, monitoring of power and cooling and failover solutions in place to ensure maximum uptime. Speed of deployment, cost-effective scalability, and ease of management with limited resources are also main concerns.

Another challenging area is the merge of information technology with operational technology. The networking of the first type is established around the open-source and Internet protocols and international standards. The second, an organization's industrial controller and automation infrastructure, is characterized by proprietary standards and non-cross functional equipment. This includes the hardware and the software that is used to control and monitor the manufacturing equipment and processes, and most communication is accomplished between machines. The two groups have separate technology stacks, network architectures, protocols, standards, governance models, and organizations. By converging them in one solution, organizations can create better products, achieve cost and risk reductions, and improve performance, flexibility and efficiency.

Recommendation 4 ***(related to the Cloud service market orientation):***

Support the ***evolution of the Cloud services towards diversification,*** *ensuring special features like those sustaining user mobility, user as service provider, service composition, personal data service configurability, or speciality Cloud services.*

This recommendation is based on the following ***visions****:*

The future of cloud service is related service consumer mobility. The flexibility demanded by the mobile workforce is a key reason for the rise of the Cloud computing and 'Bring your own device' (BYOD) is highly relevant in the world of Cloud computing. End users are using mobile devices to put their data into personal cloud services for storage, streaming or syncing. Therefore the proper means need to be put in place to integrate personal Cloud services for the companies employees in a BYOD environment.

Through connecting cloud services, the companies should be able to effectively manage their assets. However, a close watch on which providers offer cost-efficient or reliable services is needed. The ability to build agile hybrid clouds will trickle down to the mid-tier.

Social tools will bring increased collaboration to the cloud. As people become more accustomed to these tools, customers use them to communicate with fellows. The face of cloud is expected to change into a solution which will be delivered as a service with end users behaving more like service providers as a result.

Data generation activities among users and the need to access this data from anywhere using any device has propelled the demand for personal cloud for real-time data access and its sharing. Additionally, the growing need to create a backup of critical data, along with provisions for data recovery and planning for digital contingency, is boosting the personal cloud adoption consideration among enterprises as well as individuals. The personal cloud providers are expected to focus on business users along with offering superior value added services for users, supplemented by the integration of wearable technologies. The users will pay to easily deploy configurations and services they are acquainted to use. The exploitation of containers in embedded and resource constrained devices gives the opportunity to encapsulate user preferences, applications and services in virtual environments and to deploy them through the Cloud in one or many personal devices.

Specialty Clouds will raise through their ability to avoid specialized hardware acquisition costs and outdated equipment. E.g. video editing will create expanding markets for large instances and specialized GPU clouds. Running high-end graphics applications is requiring today substantial hardware infrastructure investment; newly cloud-based graphics technologies are showing that end users can run high-end graphic design applications with a simple web browser.