Concertation Report
Novel approaches and technologies for Cloud resource and service management (NATRES)

Think Cloud Services for Government, Business & Research

www.cloudwatchhub.eu|info@cloudwatchhub.eu |@cloudwatchhub

1 Identifying topics for the Horizon 2020 ICT Work Programme 2018 – 2020

Chaired by Francisco Medeiros, Deputy Head of Unit E2 Software, Services, Cloud computing

Presentation

Concertation meetings have been key in shaping Work Programmes for the European Commission. The CloudWATCH concertation meeting in September 2014 saw participants recommend themes for the LEIT 2016-2017 Programme. From this, two calls were closed just days before the 2016 Concertation meeting: ICT-06-2016: Cloud Computing, and ICT-10-2016: Software Technologies. Both were heavily subscribed and it is expected that over 20 new projects will be funded and commence later in 2016, including international cooperation projects.

The European Commission has initiated the preparation of the 2018-2020 Work Programme (WP) with a scoping paper to be prepared in summer 2016 with discussions with member states taking place between July and October 2016. Public consultation may occur in Q3 of 2016 while the WP will be drafted between November 2016 to June 2017. The WP is likely be adopted in 4Q 2017.

The four EC themed clusters announced at the 2015 Concertation meeting have a key role in supporting this process by providing white papers which include a set of recommendations for research challenges to be addressed by the 2018-2020 WP. These were presented and are outlined in the following sections.

1.1 Novel approaches and technologies for Cloud resource and service management (NATRES)

Dana Petcu, West University of Timisoara & Cluster Chair

Cluster website | **Cluster white paper** | **Presentation**

NATRES cluster is 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 include the integration of special devices from high performance computing to mobile devices, design of decentralised service-oriented systems, improvement of the virtualization technologies, overcoming portability and interoperability issues, or automation of organisation and management of back-end resources. Cloud-based applications from the fields of Internet-of-Things and Big Data are also expected to challenge the new services.

Twenty actions partially funded by European Commission in FP7/CIP/H2020-ICT programmes are actively participating to NATRES activities. NATRES results until now are reflected in (1) a map-of-challenges; (2) a white paper; (3) recommendations for the next work programme; (4) common disseminations actions in form of workshops, scientific papers, exhibition booths.

The cluster members are confident that resource and service management remain a challenging topic to be further addressed in the future ICT programmes as many related problems have not yet found a proper solution. Examples of such problems are the seamless integration of Cloud resources and service into application development tools or with other distributed or specialized systems, the full automation of datacentre services to reach multiple objectives like energy efficiency and high level of service quality service, or diversification of the Cloud services.

Cluster members

AppHub, ARCADIA, BEACON, CloudLighthning, CloudSpaces, ClouT, CloudWave, DICE, ENTICE, iKaaS, INPUT, IOStack, Mikelangelo, Mobile Cloud Networking, Mo-bizz, MODAClouds, MUSA, RAPID, SPECS, SWITCH.

The following tables provide an overview of current challenges identified by the cluster so far[1] and identified topics & recommendations for the 2018-2020 Work Programme.

[1] details at: https://eucloudclusters.files.wordpress.com/2015/05/map-of-challenges2.pdf

Table 1 Overview of current challenges identified by the cluster so far

Challenge	Description	Projects addressing issue
Deployment and management of resources: in a decentralised, autonomous way	Improving the efficiency of deployed resources; deploying of resources efficiently - use of low overhead, very efficient virtual technology stack.	-**iKaaS** is developing distributed autonomic resource (cloud, IoT) allocation mechanisms, taking into account service needs and resource capabilities and availability. -The Dynamic Real-time Infrastructure Planner subsystem in **SWITCH** prepares the execution of the applications by defining an optimal virtual runtime environment from one or more Cloud providers, and deploying the platform required by the application. -**ClouT** is developing a mechanism to manage several different kind of distributed sensors and actuators (legacy, IoT based, virtualized and “sensorized” devices). The received data are treated as real-time data for some applications and are stored in a cloud storage to be used as historical data. -**SPECS** provides solution to deploy services in Cloud according to SLA life cycle. It offers tools that automatically enforce and monitor security properties -**INPUT** is moving cloud services closer to mobile end-users and smart devices in an autonomous, energy-efficient, and dynamically fashion, in order both to avoid pointless network infrastructure and data centre overloading, and to provide lower latency to services; -**RAPID** will develop a secure peer-to-peer model where almost any device can operate as an accelerated entity and/or as an accelerator serving other less powerful devices.
Software defined execution models	Software defined networking; Software defined data center	-At the core idea of the **SWITCH** environment, a new development and execution model, an application-infrastructure co-programming and control model, will be developed for time-critical Cloud applications. The new model brings together the application composition, execution environment customisation, network programmability and runtime control, which are normally treated in separated processes, into one optimisation loop based on the time critical requirements. -**INPUT** is enabling personal and federated cloud services to natively and directly integrate themselves with extended software defined networking and network function virtualization technologies close to end-user SDs in order to provide new service models. One of the **SPECS** case studies focuses on adoption of SLA in ngDC (next generation Data Centers).
Data storage	Infrastructure service	-**ClouT** will provide a CDMI Cloud Storage to store sensor data and binary data objects (such as images or videos). The storage will be fully compliant with CDMI 1.1.1 specification and will provide scalability and elasticity to store a virtually unlimited amount of data and to manage sudden bursts. -**SPECS** provides a mechanism for End-to-End Encryption, offering even metrics to be used to grant the security in an SLA. SPECS provides an application for Virtualized Cloud Data storage protected with SLA.
Resilience and scale	Optimality; Large-scale experiments	-The ASAP subsystem in **SWITCH** autonomously manipulates the application and runtime environment to maintain optimal system level performance against time critical constraints. -**ClouT** manages and storages a virtually unlimited amount of data that will be available to citizens especially in case of emergency (e.g. natural disaster). The availability in case of emergency is one of the key use cases of ClouT. The team is working on an intercontinental demo involving the four pilot cities to show how data will be able to be shared, computed and used in real time across two continents.

Service Level Agreements	SLAs for quality critical applications and novel negotiation mechanisms	-The **SLA-Ready** Common Reference Model (CRM 1st iteration, June 2016) describes, promotes and supports the uptake of cloud service level agreements, by providing a common understanding of SLAs for cloud services. It provides a common understanding of SLAs for cloud services, integrating: SLA components like terminology; SLA attributes; Service Level Objectives (SLOs); guidelines; best practices. -The Dynamic Real-time Infrastructure Planner (DRIP) subsystem in **SWITCH** prepares the execution of the applications by creating a Service Level Agreement with the resource provider(s), and deploying the platform required by the application. **-SPECS** proposes an advanced Security SLA Model, as an extension of WS-Agreement, that enable to concretely represent security in SLAs. Moreover SPECS offers a Security Metric Catalogue, built on top of existing standard and tools offered at the state of the art. SPECS offers a Security negotiation module, devoted to negotiate SLAs according to end users needs. **-MUSA** studies how to offer Security SLA for multi-cloud applications. **-RAPID** builds tools that monitor offloaded performance of tasks when executed on the cloud. The SLA model will provide the mean to access to the monitoring information related to the agreed QoS aspects, in such a way it will be possible to determine whether SLAs are being fulfilled or not. It will be able to deal with several monitoring platforms, since it allows for customization through a system of plug-ins.
New models of services	Novel composition model for infrastructures: application aware	-At the core idea of the **SWITCH** environment, a new development and execution model, an application-infrastructure co-programming and control model, will be developed for time-critical Cloud applications. The new model brings together the application composition, execution environment customisation, and runtime control, which are normally treated in separated processes, into one optimisation loop based on the time critical requirements.
Quality assurance	QoS and QoE	-In **iKaaS**, service quality related metrics have been identified and are being considered by the service provisioning mechanisms The Dynamic Real-time Infrastructure Planner (DRIP) subsystem in SWITCH prepares the execution of the applications by semantic modelling and linking of different QoS/ QoE attributes and defining an optimal virtual runtime environment from one or more Cloud providers. -The **INPUT** framework exploits the monitoring of network and cloud QoS, and the perceived QoE of end-customer in order to enable a smart and energy-efficient management of resources. **-RAPID** take into account QoS/QoE attributes in order to select the appropriate accelerators per client. The QoS attributes of a task/application are instructed by the developer at implementation time. The SLAM monitors if these QoS attributes are satisfied, and notify the user if this is not the case.
Virtualization technology	Performance increase in virtualised IO	**-RAPID** uses OpenStack platform with KVM hypervisor on the cloud side to host different flavors of virtual machines. Each virtual machine will accommodate one acceleration server.

Table 2 Identified topics & recommendations

Topic	Description	Recommendation and benefit to future cloud market
Cloud middleware	The software needs to be described in abstracted space and needs to be practically hardware agnostic. The recent Lambda architecture, in which code is executed in response to an event, ensures 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. The application should be able to form automatic and fleeting associations of hardware and software resources according to their needs. Application developers should not 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.	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
Datacenter services	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. The datacentres need to be like biological organisms: having different states, growing and shrinking according to workloads, automatically corrected and changed. Low-power processors will be able to treat many workloads in the highly automated datacentres to support massively federated, scalable software architecture. Technologies currently limited to supercomputing will make it into the mainstream.	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 ecosystems 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.
Cloud model evolution	Inverse cloud models, like machine-to-machine computing or geo-distributed cloud, are alternatives to the bandwidth-intensive Cloud approach. However, Fog nodes for example, 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	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.

	on features such as out-of-band access, automatic detection and recovery from outages, cellular connectivity, or high-level security.	
Cloud service market orientation	The future of cloud service is related service consumer mobility. 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. Specialty Clouds will raise through their ability to avoid specialized hardware acquisition costs and outdated equipment.	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.

Conclusion and next actions for the Cluster

According to the general opinion of the NATRES cluster members and the result of online questionnaire, the evolution of the infrastructure-as-a-service model and market in the future will follow four main directions: significant changes in the support of data application development, datacenters reorganisation, model evolution towards smooth integration with other models, market orientation towards specialization and personalization.

NATRES cluster intends to support the early development of the new concepts in these four main directions by creating the framework for collaborations between its cluster members. In parallel with the challenges and recommendations, NATRES has identify the gaps between the user needs and the available technologies that will not be covered by the current cluster actions. These can be the subjects for new collaborations between the cluster members and beyond in the frame of the incoming actions.

The collaborative actions that are foreseen for the near future are related to exchanges of best practices as well as common dissemination actions in form of scientific papers or white papers, workshops or trainings organisation, journal special issues.

The list of incoming activities as well as other details about the former activities of NATRES cluster (including their results) can be found at: https://eucloudclusters.wordpress.com/new-approaches-for-infrastructure-services/