Inter-cloud Challenges, Expectations and Issues Cluster Position Paper

Initial Research Roadmap and Project’s Classification

Participant projects and representatives

Name	Affiliation	Project
Beniamino Di Martino	Second University of Naples	mOSAIC
Matteo Biancani	Interoute SpA	CYCLONE
Jose Aznar	I2CAT	CYCLONE
Domenico Gallico	Interoute SpA	CYCLONE
Ana Juan Ferrer (Editor)	ATOS	ASCETiC
Karim Djemame	University of Leeds	ASCETiC
Elisabetta Di Nitto	Politecnico di Milano	MODAClouds
Gabor Kecskemeti	MTA SZTAKI	ENTICE
Zhiming Zhao	University of Amsterdam	SWITCH
Francesco D’Andria	ATOS	SeaClouds
Robert Woitsch	BOC	CloudSocket
Kyriakos Kritikos	FORTH	CloudSocket
Peter Deussen	Fraunhofer	AppHub
Philippe Massonet	CETIC	BEACON
Massimo Villari	UNIME	BEACON
Jose Costa	AALTO	SSICLOPS
Emil Kowalczyk	ORANGE	SSICLOPS
John P. Morrison	UCC	CLOUDLIGHTING

Table of Contents

1. Introduction

Inter-cloud Challenges, Expectations and Issues Cluster goal it is to create a critical mass of projects addressing the topic of multi-cloud and inter-cloud so to share experiences, collaborate on approaches and discuss challenges for adoption and future research.

The simultaneous or serial use of services from diverse heterogeneous clouds is a challenge in order to further develop the Cloud market in Europe. While it presents a series of issues with regards to interoperability among heterogeneous cloud typologies, private and public clouds, services' comparability, portability, migration, networking, … It also offers innovative market opportunities in order to avoid vendor lock-in and for the development of new roles in the cloud market related to hybrid cloud models.

Today a number of research projects are analysing this problem from diverse perspectives and focusing on specific parts of these and other identified challenges. Within this cluster we aim to offer a forum in which to collaborate in order to elaborate a broad EU perspective of the Inter-cloud question, so to establish means of project's collaboration in two ways:

Technical	Dissemination & Exploitation
– Technical collaboration across projects – Collaboration on scientific papers – Share/Develop best practices to approach technical collaboration with Open Source Communities – Development of Future Research roadmaps	– Organisation of joint workshops with both academic and industry focus – Support to individual projects by organising project innovation management events – Approach to innovation bodies – Organisation of exploitation workshops, so to bring together researchers, technology. – Whitepapers on specific topics

The following projects are currently involved in this initiative:

Project	Name	URL
CloudSocket	Business and IT-Cloud Alignment using a Smart Socket	www.cloudsocket.eu
BEACON	Enabling Federated Cloud Networking	www.beacon-project.eu
SSICLOPS	Scalable and Secure Infrastructures for Cloud Operations	www.ssiclops.eu
ENTICE	Decentralised repositories for transparent and efficient virtual machine operations	www.entice-project.eu
CYCLONE	Complete Dynamic Multi-cloud Application Management	www.cyclone-project.eu
CLOUDLIGHTNING	Self-organising, Self-managing heterogeneous cloud	www.cloudlightning.eu
AppHub	The European Open Source Market Place	www.apphub.eu.com
ASCETiC	Adapting Service lifecycle towards Efficient Clouds	www.ascetic-project.eu
ModaClouds	Model-Driven Approach for design and execution of applications on multiple Clouds	www.modaclouds.eu
PaaSage	Model-based Cloud Platform Upperware	www.paasage.eu
SeaClouds	Agility after the deployment	www.seaclouds-project.eu
mOSAIC	Open-Source API and Platform for Multiple Clouds	www.mosaic-cloud.eu
SWITCH	Software Workbench for Interactive, Time Critical and Highly self-adaptive Cloud applications	www.switchproject.eu

This position paper establishes the first step towards the envisioned collaboration activities among these projects in two main directions:

Section 2, Inter-Cloud Research Areas and Challenges, provides an initial research roadmap for Inter-Cloud computing development areas. First, it identifies a Cluster's vision of Inter-Cloud topics development by 2020, as well as, it presents and prioritises both research areas and specific research challenges in order to make the provided vision, reality. Finally, a vision of prioritization of this research areas and challenges is provided. The objective of this work it is to serve as consultation process for work programme H2020 in the areas of Cloud Computing and Software related to Inter-Cloud and aims to reveal inputs brought from existing projects and collaborating institutions in Inter-Cloud cluster related projects.

In addition analysis of existing research projects in the Inter-cloud area it is provided in Section 3, Analysis of EU funded research projects. This classifies existing project's work in relation to Inter-operability approach, Topics and Scenarios addressed, Standards used and enabled Use Cases allows the Inter-Cloud cluster to have a solid ground in which to so to facilitate deeper technical collaborations among projects, as well as to link, proposed Research Areas for future research with on-going works.

2. Inter-Cloud Research Areas and Challenges

Multi-Cloud is defined as the serial or simultaneous use of services from diverse providers to execute an application [1]. At business level, Hybrid Cloud is the term commonly used, Gartner [3] defines hybrid Cloud as the coordinated use of cloud services across isolation and provider boundaries among public, private and community service providers, or between internal and external cloud services. A number of scenarios demonstrate these serial or simultaneous interactions among hybrid heterogeneous private and public clouds and across all cloud layers (IaaS/PaaS/SaaS)[4].

Cloud bursting is the simplest and most common hybrid/multi-cloud cloud model scenario, in which an application that is executing in a private cloud bursts into a public cloud when the demand for computing capacity spikes. The advantage of such a hybrid cloud deployment from a cloud user's perspective is that an organization only pays for extra compute resources only when they are needed.

In a federated cloud scenario, a cloud provider sub-contracts capacity from other providers as well as offer spare capacity to a federation of cloud providers. Parts of a service are placed on remote providers for improved elasticity and fault tolerance, but the initial cloud provider is solely responsible for guaranteeing the agreed upon SLA. The federated cloud scenario is related to community cloud set-ups, or from a commercial perspective, for cloud providers that own multiple cloud islands in diverse regions, in order to balance workload among them.

In a multi-provider scenario, the user, or a broker acting on behalf of the user, is responsible for management of multi-cloud provisioning of the services. Access to this functionality can be provided either directly or thought by a cloud marketplace in order to hide management complexity. The user, or an acting-broker, contacts all possible Cloud providers, negotiates terms of use, deploys services, monitors their operation, and potentially migrates services (or parts thereof) from misbehaving cloud providers. Cloud providers are managed independently and placement on different providers is treated as multiple instances of deployment.

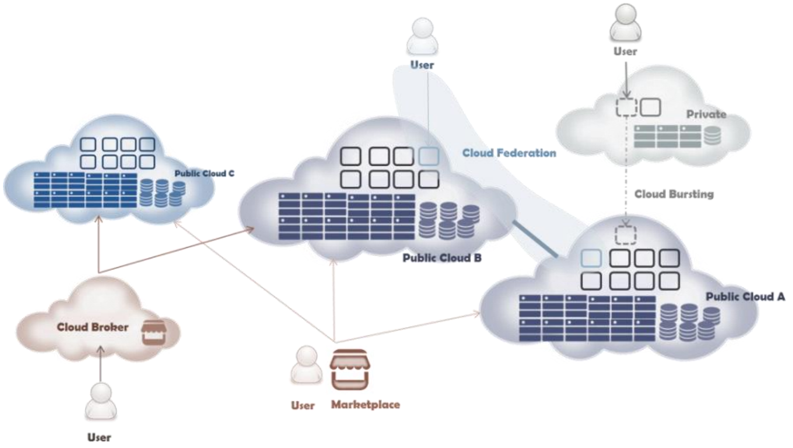

These multi-cloud scenarios are classified according to two main criteria: decision-actor and decision-time. Decision-actor determines whether the decision to use diverse clouds is taken by a cloud provider, a cloud broker or cloud user. The decision-time classifies scenarios considering if they happen at time of deploying the application or once it is in execution. The following table represents all the different interactions.

	Decision Actor		
Decision time	Cloud User	Cloud Broker	Cloud Provider Private or Public
Deployment	Multi-Cloud Cloud Brokerage Cloud Marketplace	Multi-Cloud Cloud Brokerage Cloud Marketplace	Federated Cloud
Execution			Cloud Bursting

There are a number of motivations for embracing multi-cloud set-ups both from a provider and customer's perspectives.

Provider perspective

- **Scalability and wide resource availability -** Although one of the mostly used slogans for selling Cloud computing is the argument of infinite capacity, the reality is that for any cloud computing set-up (private or public) the number or available resources are countable and finite. Often cloud providers overprovision so to keep the infinite capacity illusion and enable customers to dynamically scale their workloads, but this is not sustainable overtime neither optimal at cost level. For any cloud set-up in a private or public cloud offering the capacity limit is one day achieved. Cloud federation and busting models are the mechanism to handle peak-loads in small and medium cloud set-ups in order to acquire additional capacity only when needed.
- **Cost efficiency and energy savings**- A direct consequence of cloud providers not overprovisioning it is that costs associated to idle capacity are therefore avoided. Multi-Cloud allows, as typically specified for Cloud consumption, to reduce costs are both at level of capex and opex. Capital investments due to the fact that there is no longer the need to purchase hardware just to absorb demand peaks to get it from associated providers; but also opex, reducing overall management costs and energy consumption, for not having cloud resources underutilized.

Customer perspective

- **Avoid vendor lock-in-** Concerns about vendor lock-in are one of the major obstacles for wider cloud adoption. In current cloud status, customers are often locked to a specific cloud vendor product or service, and easily transition to a competitor does not exist. Lack of interoperability and portability spans the complete Cloud stack, embracing data, applications and infrastructure. Development of a Cloud market that considers utilizing resources from multiple providers in a transparent, interoperable, and architecture independent manner can help Cloud users to overpass existing vendor lock-in fears and develop a cloud market in which freedom of choice prevails.
- **Distribution across geographies for reducing latency, address legal constraints and enable high availability -** The multi-layered nature of Clouds present in its current status concerns for users with regards to regulatory context. Existing world-wide established providers address this issue, by offering diverse regions with limited level of automation among these. This mechanism it is also offered to support high availability. Beyond these, increased automation among diverse cloud offerings in different geographies can satisfy increasing demands for user businesses to act at a global scale, fulfilling specific applicable regulations, automating high availability across clouds while addressing needs spread service consumers.

These motivations for multi-cloud hybrid models in order to become a reality in a mature Cloud market need of the realization of a series of new technological developments together with advances in standardization efforts addressed in the following vision for 2020, research areas and challenges.

Inter-Cloud vision for 2020

Despite the achieved advances and commercial uptake, Cloud technologies and models have yet to reach their full potential. Many Cloud capabilities need still to be further developed and researched, so to allow their exploitation into a full degree. All along the Cloud stack (SaaS, PaaS, IaaS) commercial product developments today are based into proprietary solutions that drive to a vendor lock-in situation for the existing adopters. In this context, the realisation of multi-clouds is materialised though internal clouds and interactions between public-private Clouds which is hardly automatized and, in any case, automatic. In addition, security, trust and legal compliance issues still act as barriers for a wider uptake. Whilst more developed Inter-cloud scenarios, such as Cloud Bursting, Cloud aggregation and Cloud brokerage exist theoretically, real implementations marginally exist and they are tailored for specific cases. To reduce the effort and time associated to the adoption of cloud, developers need to be able to develop an application regardless of where it is released, structuring and building it in a vendor agnostic way so that it is possible to deploy on the provider that best fits the requirements at the moment thus realizing the “develop once deploy everywhere” paradigm.

Today Cloud Computing market is still far from adopting an open and competitive model in which cloud resources act like in conventional markets. Lack of interoperability and adopted standards together with intricate regulatory context, inflexible pricing models and not adequate SLAs are recognised as the main obstacles to Cloud adoption. However, in order to realize a full Multi-Cloud market vision additional aspects need to be developed into Inter-cloud management such as: provisioning, metering and billing, privacy, security, identity management, fine grained QoS and Service Level agreements, consideration of diversity of resources (compute, data and network).

The use of standard or agnostic interfaces for cloud services would allow the developers to migrate cloud application among cloud platforms with minimum effort. This alignment need to be achieved at all cloud levels and across different models of clouds (including local/edge clouds).

Automatic porting of existing applications and software systems (in particular legacy systems) from on premise platforms to a cloud platform need to be supported by suitable methodologies and tools to facilitate and speed up the migration.

It is solely with the full development of these novel capabilities and when real interoperability among cloud providers will become a reality through the adoption by market leaders of existing or innovative standardisation efforts, that new scenarios of business opportunities will emerge both for existing and new stakeholders in Cloud computing and EU Single Digital Market for European businesses. This will allow for the exploitation of multi-cloud models to their full potential, enabling the shift from a product centric provision into a service-oriented economy in a rich Cloud ecosystem. In this ecosystem customer-driven dynamic composition of cloud services will allow to provide customer-tailored complex services, creating economic value from the interconnection of diverse and heterogeneous service providers that jointly contribute to an integrated solution that meets individual customer needs. Targeting a Cloud Ecosystem where freedom of choice prevails for the customer and where all cloud stakeholders (Cloud Service Providers, Software Vendors and Telecom operators) leverage and multiply the benefits of each other taking into account the EU Single Digital Market for European businesses. Automatically discover and compose cloud service at different levels (e.g., business process, software, infrastructure) in order to satisfy application or business requirements, will enable not only the fast development of applications and business processes for the cloud, but also their runtime adaptation, when the respective need arises, thus catering for dynamism.

The advent of these models will bring significant challenges to customers, in which IT provision will deploy a hybrid IT model going beyond conventional approaches. For several decades, customers and providers have relied on stable delivery approaches, customized architectures and solutions and traditional commercial models, with outsourcing at the heart of their offerings. In a multi-cloud market, enterprise IT services will not be organized neither only in-house nor completely outsourced, instead will be located at some optimal point between the two and enabling vertical market customization, realized by the ability to assemble, consume and operate cloud services taking into account operational concerns and technological challenges applicable to a specific sector. This will allow easier outsourcing of business processes that do not constitute a competitive advantage to customers and the creation of cloud markets that combine general purpose services and specific services tailored for the needs of vertical markets emerging in the context of EU Single Digital Market.

Assuming that the provisioning of the heterogeneous cloud infrastructures is in place, it becomes a challenge to monitor and react upon unexpected degradation of service quality, by identifying the source of the problem, in the specific cloud provider. The coordination and infrastructure re-configuration (possibly involving the other cloud infrastructures) is key to ensure the restoration of SLAs and guarantee the proper behaviour of applications and services.

This will happen into a general trend towards specialisation and decentralisation also affecting Cloud providers, and in general, the cloud market as a whole. Hybrid Inter Cloud models will have to go beyond private – public cloud interactions but to consider a wide range of clouds typologies and services. These have to take into account specialised sector services but also new technological developments addressing specialised set-ups for large data storage and intelligent analytics and Edge/Fog Local clouds for IoT.

These will require Multi-Cloud management techniques to overpass its current central actor, broker or customer, that controls the action to a real autonomic Inter-Cloud management layer, having the workload a keystone, so to optimise resource utilisation and find the best location in which to execute as part of a novel compute continuum, enabled by simple and transparent workloads movement across diverse cloud and resource typologies and models. Movement towards hyper distribution of computing will have to enable cloud computing continuum scenarios to consider interoperability, portability, elasticity, self-organisation, self-management and self-healing across many and heterogeneous resources in micro local clouds, private enterprise clouds, aggregated cloud models and large Cloud set-ups. Orchestration and placement problems in this context require of consideration of heterogeneity and trade-offs among consistency, reliability and performance. Novel Programming models and software engineering approaches will have to emerge to cope with hyper distribution of computing considering enhanced workload portability abilities, richer fault-tolerance verification with complexities at level of extreme scalability and parallelisation.

Overall this will be a significant breakdown for existing Inter-Cloud computing developments, which emerged as part of a centralised paradigm, to a pure Hybrid in a wider sense and multi-cloud decentralised and autonomic management model, based on decentralisation and connecting a wide diversity and variety of entities and typologies of resources requiring new developments in a novel Inter-cloud computing continuum.

Research Areas

In order to make reality the vision sketched in previous section, the following research areas have been identified by the Inter-Cloud cluster.

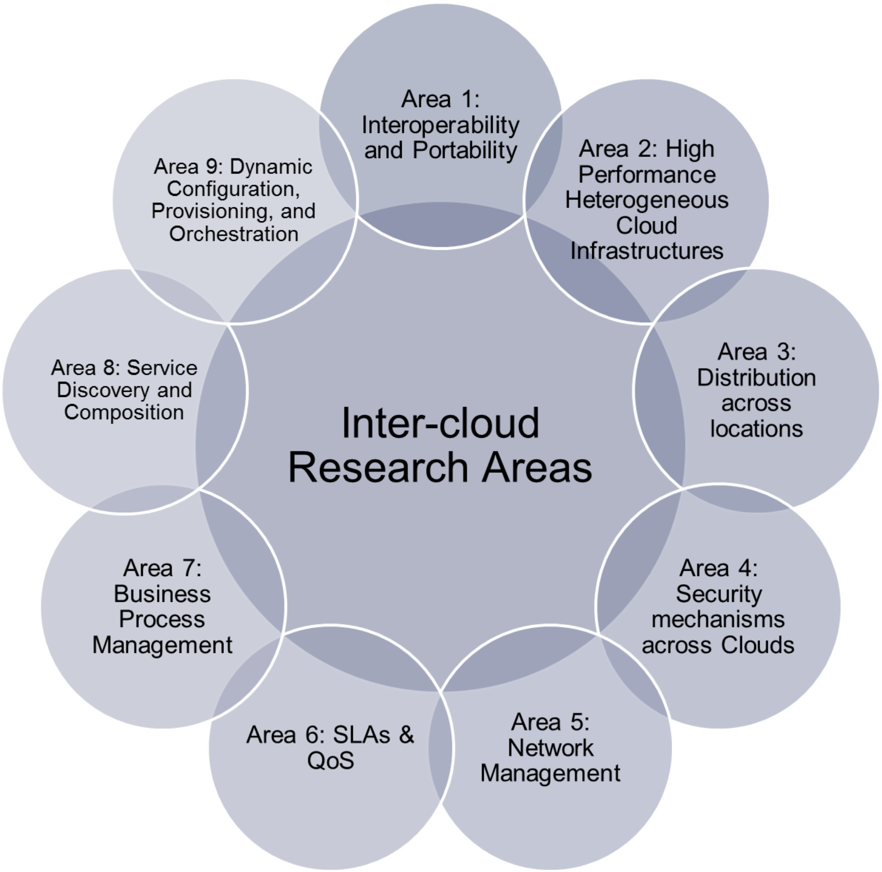

Figure 1 Inter-cloud Research Areas

Area 1: Interoperability and Portability

In today's Cloud market there are dominant vendors that stick customers onto their specific technologies of choice and shuns that customers move to different providers without significant costs and technical efforts. Nowadays a very significant number of standardization efforts are happening all across cloud stack from many and diverse standardization bodies. The EU reports [26] that list of around 20 relevant and active organizations in cloud computing standardization area from which around 150 associated documents, standards and specifications are currently available. Unfortunately it is still to happen that Cloud market leaders adopt any of them widely. The trigger to multi-cloud and hybrid cloud can be the market force that pushes the adoption of standards by major vendors, breaking down current vendor lock-in situation.

Moreover, in order to enable automatic establishment of chains of contractual relationships across multiple and heterogeneous cloud providers, it will be necessary to analyse and extend when necessary, existing formalisms to describe Cloud service offers in order to enable service comparability or "fungibility". This is, the capacity of services to be compared or are capable of substitution. Fungibility will enable to act in two aspects: firstly, by enabling further reusability of atomic services by participating in multiple complex cloud service provisions; and secondly, by enabling the optimization of cloud service compositions at operation by replacing services that do not provide the adequate QoS performance.

Traditionally the aspect of workload portability across cloud providers has mainly focused on VM portability. Today, VM portability is hindered to be focused on single provider due to the need of shared to the storage and network services in the same local area network. The consideration of multiple cloud environments and innovative virtualization mechanisms opens the door to very interesting capabilities and new developments. These developments will enable higher efficiency, higher performance and easier application encapsulation allowing better abstraction from resources. These benefits will be extended to the Orchestration and Contextualization areas, replacing much of what it is currently decided by orchestration means in the cloud provider at deployment time. Differently, these instructions can soon be part inside the virtualization format, enabling multi-cloud highly automated systems to accept those instructions and building cloud applications out of them, independently of cloud's technologies of choice.

Associated Future Research Challenges

Challenge 1. Develop once deploy everywhere

Developers need to be able to develop an application regardless of where it is released, structuring and building it in a vendor agnostic way so that it is possible to deploy on the provider that best fits the requirements at the moment thus realizing the "develop once deploy everywhere" paradigm.

Challenge 2. Switch services among cloud typologies and providers without efforts

The use of standard or agnostic interfaces for cloud services allows the developer to migrate cloud application among cloud platforms with minimum effort. This alignment need to be achieved at all cloud levels (IaaS, PaaS and SaaS) and across diverse models of clouds (local/edge clouds).

Challenge 3. Interoperability to cope with Cloud heterogeneity and application mobility

Management of potentially thousands/millions of small diverse devices and sensors in a fog computing set-ups combined with multi-cloud approaches will require of new management styles and interoperability mechanisms in cloud architectures so to cope with heterogeneity of underlying resources and enable transparent workload mobility across micro local clouds, private enterprise clouds, aggregated cloud models and large Cloud set-ups. This will also need to consider speed of deployment and cost-effective scalability across clouds.

Challenge 4. Automatic migration of in house application to the Cloud and across cloud typologies.

Automatic porting of existing applications and software systems (in particular legacy systems) from on premise platforms to a cloud platform need to be supported by suitable methodologies and tools to facilitate and speed up the migration and adequate QoS levels . Further advances, will have to consider transparent migration of applications or application components, across Edge local clouds and network distributed clouds.

A big boost to make feasible these desiderate features could be represented by semantic technologies that enables the definition of an abstraction layer above existing concepts thus allowing the interoperability of data and services, as detailed in Challenge 6.

Challenge 5. Extended Workload Portability

Support for new technologies, such as containers, in order to provide higher efficiency and performance, as well as improved application encapsulation and abstraction from resources while porting applications across diverse and heterogeneous clouds.

Challenge 6. Universal Semantic Service Description

There is a need to develop a universal service description language which can rely on semantics and which will be able to cover both functional and non-functional aspects. This language should be able to describe cloud services in different levels of abstraction. Through the use of this language, the comparability between cloud offerings can be established as it will guarantee that common description terms are used. Moreover, the accuracy in cloud service discovery will be enhanced. In this way, derived cloud service compositions will be more robust and suitable when derived from the respective composition tools. In addition, service substitution could be more effectively applied in cases, e.g., one cloud service is under-performing. To move from structural or text-based cloud service descriptions to semantic ones, it is imperative that the language is complemented with semantic alignment tools, which are able to map terms in the structural or text-based descriptions into ontology concepts, as well as respective transformation tools.

Prioritisation

	Challenge
5	Challenge 2. Switch services among cloud typologies and providers without efforts
4	Challenge 1. Develop once deploy everywhere
4	Challenge 3. Interoperability to cope with Cloud heterogeneity and application mobility
4	Challenge 6. Universal Semantic Service Description
3	Challenge 4. Automatic migration of in house application to the Cloud and across cloud typologies.
2	Challenge 5. Extended Workload Portability

Area 2: High Performance Heterogeneous Cloud Infrastructures

The possibility to allow users to aggregate cloud resources from both private and public providers to build a cloud platform that is tailored to their application's needs may add great value while deploying applications in inter-Cloud scenarios. This tailor-made platform may include both high-performance computing and high-throughput computing resources. The high-level abstraction of cloud- and application-level services will permit users to take advantage of heterogeneous cloud resources with minimal effort.

Associated Future Research Challenges

Challenge 1. Enable with inter-Cloud Service Provider connectivity

While aggregating the resources of both local, public and private cloud platforms, it becomes a challenge to connect the resources that have been requested to them and enable to the application level with the means to configure certain networking aspects required by their deployed applications or services (i.e. IP addressing, application firewalling options, load balancing mechanisms, etc.). Not all the cloud platforms expose these type of networking information, so that it becomes a challenge to identify (1) what each Cloud provider enables to handle and (2) aggregate such options to be offered to the application level.

Challenge 2. Monitor and guarantee inter-cloud infrastructure SLAs performance

Previous challenge was related to the provisioning of heterogeneous cloud infrastructures. Assuming this is in place, it becomes a challenge to monitor and react upon unexpected degradation of service quality, by identifying the source of the problem, in the specific cloud provider. The coordination and infrastructure re-configuration (possibly involving the other cloud infrastructures) is key to ensure the restoration of SLAs and guarantee the proper behaviour of applications and services.

Challenge 3. Dynamic workload balancing in multi-cloud context

When considering a multi-cloud context, the balancing of workload needs to be concerted both among the different platforms involved and within each of the specific Cloud infrastructure: thus, it is necessary to consider two levels of granularity for load balancing. At a higher level, the main challenge regards the effective distribution of the workload among the heterogeneous Cloud platforms provisioned, which needs to take in consideration the SLAs agreements (the monitoring of which is covered by Challenge 2), the capabilities of each considered platform and, most importantly, the dynamic and ever changing requirements of the running distributed applications. The feasibility (and actual implementation) of a possible workload re-allocation has to be evaluated dynamically: in this case, apart from the time\network resource costs involved in such operation, it is necessary to take into account potential data incompatibilities or services' interfaces inconsistencies which may exist among different platforms. At platform level, the workload is further distributed among the existing services instances, whilst new ones need to be created as needed for scalability purposes. Here the challenge is to identify the optimum number of virtual instances to create and destroy according to the workload to distribute, taking in consideration the time needed to perform such operations.

Challenge 4. New languages to express overall high performance including storage, compute, network

There is will provide new languages and mechanisms to co-ordinately select and orchestrate resources across connected heterogeneous cloud computing set-ups, including micro local clouds, private enterprise clouds, aggregated cloud models and large Cloud set-ups. These have to include computing and storage needs but also use of sensor and network infrastructures to support both rapid service deployment and service migration for dynamic operation of workloads in order to achieve high-performance and sustainability including eco-efficiency in multi-cloud environments.

Challenge 5. Dependability and reliability between Cloud providers and consumers

Cloud model depends on a high degree of dependability and reliability between Cloud providers and consumers; relying on the trustworthy context established between both parties. The open, dynamic and self-service nature of Cloud makes the relationships between them highly dynamic, allowing entities to join and leave frequently. This scenario is foreseen to evolve as the interoperability between providers is guaranteed, moving to highly non-locked market of customers and providers where relationships are established in an on-off basis based on the mutual agreement. In addition, as the Cloud market expands, the degree of anonymity between these entities is going to be incremented. This together with the fact that, decisions to select specific providers will be taken in more automated manners, make evident the pressing necessity of mechanisms to establish a circle of trust between Cloud services consumers' and providers' in a multi and hybrid cloud ecosystem.

Prioritisation

	Challenge
5	Challenge 2. Monitor and guarantee inter-cloud infrastructure SLAs performance
4	Challenge 4. New languages to express overall high performance including storage, compute, network
3	Challenge 5. Dependability and reliability between Cloud providers and consumers
2	Challenge 3. Dynamic workload balancing in multi-cloud context
1	Challenge 1. Enable with inter-Cloud Service Provider connectivity

Area 3: Distribution across locations for reducing latency, address legal constraints and enable high availability

In a world where digital businesses increasingly act in a global scale and not specific for any geography, it is really important to provide customers with cloud services that have mechanisms to dynamically coordinate load distribution across locations and cloud typologies. This need is threefold. First, getting workloads closer to where demand is happening avoids unnecessary latencies. But also, it allows fitting national regulations applicable to specific geographies in many situations in which customers have specific restrictions about the legal boundaries in which their data and application can be hosted. Not to forget, it is the need to automate high availability across clouds (typically private to public but local clouds and across cloud locations). Using multiple clouds simultaneously is the only solution for satisfying the requirements of the geographically dispersed service consumers. For this, the developments of multi-cloud computing environments are necessary to facilitate provisioning of application services that are fault-tolerant and satisfy globally SLA targets under variable load, resource, and network conditions.

Associated Future Research Challenges

Challenge 1. Scalability across clouds based on demand

Existing scalability mechanisms in Cloud computing typically take into account a single Cloud installation and commonly it is based on application components replication based on monitoring rules. Richer scalability mechanisms, both at vertical and horizontal levels, are required in order to dynamically coordinate load distribution across locations and cloud typologies. These have to go beyond simple reactive approaches and consider predictive analysis, and additional demand related parameters such as demand location.

Challenge 2. Cross-cloud VM/container image distribution SLAs

One of the biggest latencies a cloud user faces is the VM instantiation latency. To allow customer oriented flexible instantiation operations, VM images must be distributed to the image repositories of several cloud providers while considering storage costs. To allow the easy definition of the instantiation latencies acceptable at various geographic locations, users should be able to define their needs in an image distribution specific SLA. Such SLAs should be used while corresponding with special, cloud independent image distribution entities/components that are capable to organize image storage contents on behalf of their users while keeping their storage costs minimal.

Challenge 3. Novel High Availability mechanism across hybrid cloud models

A common usage of hybrid Cloud computing models today is in scenarios in which claims to provide high availability between private and public clouds enabling applications to remain operable (from this time, available) with one, or even multiple, application components failing. Cloud developers need to design and develop Cloud applications considering these failure possibilities, typically by enabled by usage of centralized common data-stores or load-balancers. However, this approach does not fit or apply to all kinds of workloads, and what it is more, add an additional and specific model for cloud applications. New mechanisms for supporting high availability across diverse, and hybrid in the wide sense, are required also providing transparency to the users.

Challenge 4. Legal aspects

Cloud hybrid scenarios are defined by continuous flow of data. This brings uncertainty with regards to applicable various data protection legislations and regulation that transcend national borders, therefore making complex compliance with European legislations. In addition to these, multi-cloud users need of safeguarding data's privacy when handling confidential and private customer data. There is the need of providing adequate mechanisms to manage data protection and privacy in Cloud environments. So to ensure data privacy, it is also important to contribute to standards and policies created by industry organizations, commercial enterprises and governments.

Prioritisation

	Challenge
2	Challenge 3. Novel High Availability mechanism across hybrid cloud models
2	Challenge 4. Legal aspects
2	Challenge 1. Scalability across clouds based on demand
1	Challenge 2. Cross-cloud VM/container image distribution SLAs

Area 4: Security mechanisms across Clouds

Cloud security issues are prime concerns for multi-cloud models. In order to deploy and manage a solution on multiple cloud platforms, consumers face significant challenges both due to the interface diversity and architectural differences in the diverse cloud platforms. Besides, the constant changes in security parameters enabled by dynamic multi-cloud management models would amplify current security concerns. Constant changes in the workloads enabled by migration across a wide diversity of cloud models ranging from local /fog clouds and private and public clouds, would lead to no pre-defined network topologies for applications in which resources in diverse multi-cloud models are not fixed to any geographical location.

Associated Future Research Challenges

Challenge 1. Security mechanisms for application integrity

In relation to virtual machine images (or container images), security is critical for keeping the integrity of the applications encapsulated in the images and ensuring that the contents of the images are not accessible for those who are not authorized to access them. This is especially important in the multi / inter cloud scenarios when the images could arrive to lesser security image storage systems or pass over public networks. For this reason it is essential to obfuscate secured applications and store/transfer them mostly in encrypted forms.

Challenge 2. Federated Authentication for non-Browser HTTP Applications

It is very challenging to provide federated authentication for non-browser HTTP applications, e.g., RESTful command-line tools: while there is the SAML 2.0 Enhanced Client Profile (ECP), it is unsupported by a number of tools and federations, most notably EduGain. Some tools, such as OpenID Connect, where there is a so called “Direct Access Grant”, works only for local IDPs and not for a whole federation. Of course there is also Kerberos, but it is not designed for easy federation and HTTP applications. So new mechanisms and tools are necessary going beyond of any established mechanism for this purpose.

Challenge 3. Federated Authorization Policies and Use Cases

There are quite a number of use cases, where applying XACML makes immediate sense, e.g., when huge companies want to consolidate their authorization policies and want to reuse roles and responsibilities in multiple systems. But for federated scenarios, as far as we know, there are no prominent use cases. Therefore, it is currently unclear how and why XACML can be applied in federated authorization scenarios. One example could be enabling User A to delegate access to their cloud resources to User B. But there are many other options for realizing this, e.g., OAuth 2.0. All options should be thoroughly examined and compared in order to focus on the specific needs XACML will meet within federated cloud scenarios and its benefits and drawbacks.

Challenge 4. Definition of Security and network-aware application requirements

Security needs to be considered at the design level, also providing new customizing network overlay protection. Customers in creating new applications should be able to specify the required security services considering also networking. They have to be provided by all federated Clouds to ensure that the federated cloud infrastructure meets the security requirements of the deployed federated cloud application. These security requirements must be defined during a deployment of any application also considering the design of networks.

Challenge 5. Auditability in Cloud Federated Cloud Networks

When stringing the current concept of Cloud federation so to include not only the IaaS level considered today, but also PaaS and SaaS levels together with dynamic network management and federation of IoTs additional security needs raise with regards to audit the security of software defined networks (SDN), network function virtualization (NFV) and service function chaining (SFC) in cloud networking federations. This is challenging because of the cross-cloud nature of federated cloud networks.

Prioritisation

	Challenge
5	Challenge 5. Auditability in Cloud Federated Cloud Networks
4	Challenge 1. Security mechanisms for application integrity
4	Challenge 4. Definition of Security and network-aware application requirements
2	Challenge 2. Federated Authentication for non-Browser HTTP Applications
2	Challenge 3. Federated Authorization Policies and Use Cases

Area 5: Network Management

Poor network performance is perceived as one important limiting factor in Cloud performance and identified as a main roadblock for wider cloud adoption[24][25]. It is only with the advent of virtualization and SDN that the idea of a fully scalable, end-to-end, software-based infrastructure can be realized, leveraging the features that network connectivity services may bring to Hybrid Cloud federations.

The cloud today, to some extent, represents a consolidated technology that is powering new emerging technologies like IoTs, Big Data, Open Data, along the definition of new networks SDNs and Network Function Virtualization (NFV). The current concept of Federation should be enlarged to cover not only the issues of the well-known IaaS level, but also PaaS and SaaS levels: an overlay federation is meant. In the IaaS there are confined the management of Cloud Resources (service discovery, matchmaking, VM migration, etc.) along with the federation of dynamic networks like NFV/SDN networks, even a context where it is also possible to consider the federation of IoTs (i.e. sensors and actuators, etc.). In the other two layers there are confined the federation of hybrid data analysis (example partializing the access to Data when they cannot be totally Open), identity federation (where different organization domains need to interact each other) and even the federation of IoTs but characterized by advanced services (example taking into account all services existing in Smart Cities scenarios, where a plethora of IoTs with different aims will be widely deployed). Here the new networks play an important role in connecting devices and datacentres over different network carriers. It is necessary to setup basic environments including networks where it is possible to address the above reported challenges. In particular, taking into account the involvement of heterogeneous systems along with distributed entities able to offer services and hooks is really useful to be managed in federation.

It is necessary to deal with advanced cloud infrastructures, networks, and services able to expose a rich set of APIs (standardized or to be standardized). Nowadays, Cloud infrastructures and services are looking at new DevOps paradigms; hence even networks should be designed and monitored during the DevOps Agile development and deployment.

For intra and inter cloud networks an agile setup is required along with focused actions useful for preventing misbehaviour at network level (routing optimization, security and data privacy preservation, attacks). Moreover future networks should respond more fluidly to changes in user demand at cloud level but also at the edge level (close to the IoT devices). This is the prerogative of the fog computing. The idea behind the fog computing is to bring the cloud closer to the edge and users as fog. The network represents the backbone in charge of this approach, where new hooks and multi tenancy capabilities will foster these accomplishments.

Associated Future Research Challenges

Challenge 1. Extension to Cloud Federation concept and tools

Federated Cloud Networking requires extensions to the concept, techniques and primitives of cloud networking to embrace a) federated intra- and inter-cloud networking and b) application- and service-aware virtual networking, going beyond offering it as a pure infrastructure service that is fully isolated from the cloud value chain it serves c) Heterogeneity, considering PaaS and SaaS layers, and Fog computing models, creating homogeneous overlay networks on-demand over heterogeneous clouds in multi tenancy.

Challenge 2. To guarantee new paths for optimizing transfer of data among clouds, among IoTs and clouds-IoTs

Specific efforts with regards to the definition of a network service management platform into Inter-Cloud management platform, including IoT scenarios, so to allow the dynamic allocation of network services inside and between heterogeneous clouds to improve the performance of applications that depend heavily on data access and analysis.

Challenge 3. Enablement of responding more fluidly to changes in user demand at inter-cloud level but also at the edge level

Dynamic setup of basic environments including networks, so to address the heterogeneous Cloud federation scenarios over different network carriers and considering software defined set-ups, fluidity in demand and assuring the auto-healing net paths, where multiple circuits are conceived and transparently used respect to any application requirement and needs.

Challenge 4. DevOps Agile development and deployment considering network management
Accomplishing systems designed and monitored during the DevOps Agile development and deployment phases, but even capable to create dynamic and compelling agile networks.

Prioritisation

	Challenge
4	Challenge 2. To guarantee new paths for optimizing transfer of data among clouds, among IoTs and clouds-IoTs
3	Challenge 1. Extension to Cloud Federation concept and tools
3	Challenge 3. Enablement of responding more fluidly to changes in user demand at inter-cloud level but also at the edge level
1	Challenge 4. DevOps Agile development and deployment considering network management

Area 6: SLAs & QoS

Service Level Agreements (SLAs) play a key role by being the mechanism that users have to enforce guarantees around performance, transparency, conformance and data protection when using Cloud services. While still these rich SLAs and its management techniques in single provider scenarios are not yet widely available in commercial providers, Multi-cloud environments pose higher levels of complexity to the issue. The development of multi-cloud Bursting, Brokerage or any type of multi-cloud scenarios relies in the principle that providers do not offer Cloud services directly themselves but rely on a more complex cloud ecosystem. This ecosystem enables better and more advanced services at a reduced price, but relies on third parties to provide a global service SLA. An example of this is when a provider supplying critical services with high SLA commitments chooses to migrate non-critical services to an external provider in order to ensure that the critical service receive all resources required to fulfil its SLA. In any multi-cloud scenario, the entity that acts as a mediator between the cloud user and associated cloud providers has to select services from different providers that all together best meet the user requirements, in order to offer a global service that can be produced by the compositions of several individual cloud provider services. User applications' SLA in such an environment uses services and resources from different providers, one role of the mediator is to set up and enforce a global SLA. The composition of SLAs across providers as well as the dynamic substitution of cloud services not behaving adequately at runtime is still an open issue.

Associated Future Research Challenges

Challenge 1. SLA Standard Representation

A uniform and standardised representation of SLA, agnostic with respect to different cloud providers allows specifying cloud infrastructure requirements in an agnostic way and allowing an automatic comparison among the different cloud offers coming from different providers. The language used to specify such SLA representations should be able to capture various information aspects that are required in order to support the SLA-based management of the respective services involved, including adaptation actions to be taken when SLOs are violated which could take even the form of SLA re-negotiation. The SLA specifications of such a SLA language could be complemented with the unambiguous and precise description of the quality parameters involved in the Service Level Objectives (SLOs) of this SLA. Such a description could rely on existing semantic formalisms for service quality specification. The external specification of quality terms facilitates the modeller in focusing on the main aspects of the SLA and enables the involved parties to examine such external specifications only when the respective need arises (e.g., monitoring has to be performed according to the external specification of a particular quality metric). Finally, the SLA language should be flexible enough in order to accommodate for capturing all or at least some of the possible different types of SLA composition: (a) composite SLAs which comprise SLOs at different levels and multiple signatory parties; (b) chains of dependent SLAs at different levels.

Challenge 2. Monitoring of QoS and application level monitoring

The developers need to have a provider independent monitoring system able to monitor quality parameters at different levels of abstraction that constitute key points of the signed SLAs (such as availability) in the context of SLOs. The monitoring of the system at different levels and the capturing of dependencies between these levels enables performing root cause analysis such that any particular problems at the infrastructure level can lead to automatic cloud infrastructure reconfiguration that best fits the runtime application requirements.

Challenge 3. Intelligent Broker

Cloud developers need intelligent brokering support in order to acquire the resources that best satisfy the initial application requirements and then adapt the brokering policies to the changing requirements of the runtime application.

Challenge 4. SLA-based cloud service/application management

SLAs can be used as an instrument through which cloud services and respective applications can be managed during their life-cycle. To this end, SLAs should be able to capture all appropriate information aspects relevant for this management. As such, the respective management system that will be developed can rely on the modelled information in order to appropriately realise the functionality required for each activity at the service/application life-cycle. Semantics can be employed in order to automate as much as possible these lifecycle activities. Moreover, the management system should be equipped with respective techniques, methods and tools that go beyond the current state-of-the-art in each activity such that, in the end, the cloud-based service application is managed in a holistic and effective manner across different layers in multiple clouds according to the functional and non-functional requirements posed by the respective stakeholder/end-user.

Prioritisation

	Challenge
5	Challenge 3. Intelligent Broker
4	Challenge 1. SLA Standard Representation
4	Challenge 2. Monitoring of QoS and application level monitoring
4	Challenge 4. SLA-based cloud service/application management

Area 7: Business Process Management

Currently we can identify barriers in cloud usage due to the big gap between pragmatic, legally influenced and well-defined business processes and a gigantic cloud market with numerous offerings that rarely consider the business episodes of an entrepreneur but focus on technical details. Hence, Business and IT-Cloud alignment can be realised via smart brokerage to bridge the gap between: (i) Domain specific business processes that describe the business activities of a worker, which are not executable, neither by a workflow engine within or outside the cloud; (ii) Executable business processes represented by workflows orchestrating the interaction between software applications; (iii) Cloud deployable Workflow Bundles packaged for cloud deployment consisting of all relevant deployment configurations, as well as finally (iv) Production Workflows offered in the marketplace similar to the way SaaS are supplied which can be immediately run by the user.

Various challenges exist that have to be handled in order to bridge the business to IT gap as well as leverage on the concept of Business Process as a Service (BPaaS). These challenges span: (a) enable the mapping from conceptual business process down to the level of production workflows by using semantics as well as reusing (previous) design knowledge as much as possible; (b) concrete and complete handling of instances in a multi-cloud environment in the context of a domain-specific BPaaS; (c) sustainment of guaranteed service levels for BPaaS in highly open and dynamic multi-cloud environments; (d) enable the adaptive and cross-layer provisioning of BPaaS; (e) leverage smart, semantics-based monitoring and conceptual analytics for continuous BPaaS evolution and improvement; (f) better exploitation for BPaaS providers of the variety in cloud offerings to avoid vendor lock-in as well as to select the best possible cloud services across multiple clouds to achieve desired service levels as well as maximize gains by also enabling the offering of flexible pricing models; (g) address the heterogeneity in various languages involved in the description of cloud services and workflows in the design of BPaaS.

Associated Future Research Challenges

Challenge 1. Smart business-to-IT alignment.

Starting from business requirements and respective business processes, there is a need to go down until the level of production workflows possibly running in multiple-clouds in order to support the realisation of the vision of Business Process as a Service. The mapping between the various artefacts at the different levels can be facilitated through the introduction of semantics as well as by the proper exploitation of previous design knowledge. It should also be facilitated via the addressing of the heterogeneity of the different languages that are involved in the specification of cloud services and workflows such that it will be possible to discover and exploit those cloud services which jointly, in the form of a workflow, lead to the satisfaction of the business requirements posed.

Challenge 2. Cross-layer and Scalable Multi-Cloud Workflows and BPaaS

Workflow and BPaaS paradigms need to be extended so to support flexibility, dynamicity and scalability across diverse and heterogeneous multi-tenant cloud environments. These have to consider differentiation in scope of actors and entities participating, but also the necessary self-management capabilities so to cope with specific requirements and evolution and adaptation of multi-cloud workflow itself.

Challenge 3. Cross-layer BPaaS Monitoring & Adaptation

Based on the fact that a BPaaS is a composite artefact spanning different layers, it is imperative that there are suitable mechanisms which enable its adaptive provisioning across multiple clouds. Such mechanisms should be able to monitor information at the lowest layers and propagate it up to the higher levels, in accordance to the dependencies between these layers, as well as assess which low-layer events are the root causes for the violation of high-layer requirements. Apart from detecting violations of business and technical requirements, there is a need to react on these violations on a holistic manner which respects the dependencies between the layers and does not lead to executing adaptation actions independently at each layer. Rule-based approaches could be employed for this which could focus on mapping patterns of events into complex adaptation workflows including actions specific to layers as well as novel, composite cross-layer actions.

Challenge 4. Intelligent Allocation of BPaaS across cloud levels

Driven by the business and technical requirements, there is a need to concretise the technical workflows derived from business processes. Such a concretisation needs to be performed at different cloud levels and can involve exploiting services from multiple-clouds. SaaS services are needed in order to realise the functionalities of the workflow tasks, while IaaS services are needed for the deployment of the internal workflow components. Due to the dependencies between the requirements at different levels, it is imperative that SaaS and IaaS selection is not performed in an independent but joined manner in order to guarantee that finally all requirements across all levels are satisfied. As such, there is a need for employing sophisticated optimisation algorithms, possibly exploiting different constraint solving techniques, which are able to solve a single constraint optimisation problem to enable achieving the required intelligent cross-cloud-level BPaaS allocation in the end.

Challenge 5. Smart Business Intelligence through cross-layer BPaaS Evaluation

Apart from dynamically reacting on critical situations that might be raised during BPaaS execution, it is imperative that smart business intelligence algorithms are employed such that insights on recurring problematic situations are derived which can enable optimising the performance and configuration of the affected BPaaS. Such algorithms could focus on different types of derivation which can be suitable for BPaaS optimisation: (a) detection and root-cause analysis of KPI violations which goes from the business down to the technical level; (b) identification of optimised alternatives with respect to multi-cloud BPaaS deployment which can enable achieving better service levels than the ones promised or enable the BPaaS to surpass current performance problems; (c) discovery of event patterns leading to KPI violations and transformation of such patterns to cross-layer adaptation rules which can effectively drive the adaptation behaviour of the

BPaaS; (d) mining the workflow logs of the BPaaS in order to discovery discrepancies between what has been modelled and what is evidenced during runtime, thus leading in the end to the evolution the BPaaS workflow.

Challenge 6. Flexible Cost Models

In order to guarantee the sustainability of a BPaaS and reduce the levels of risk with respect to its offering, there is a need of deriving novel cost models which are able to surpass the heterogeneity in the cost models involved in the offering of the constituting cloud services (SaaS and IaaS) as well as provide certain levels of flexibility in order to cater for a win-to-win situation for both the BPaaS Broker and Customer.

Prioritisation

	Challenge
4	Challenge 2. Cross-layer and Scalable Multi-Cloud Workflows and BPaaS
4	Challenge 1. Smart business-to-IT alignment.
3	Challenge 6. Flexible Cost Models
2	Challenge 5. Smart Business Intelligence through cross-layer BPaaS Evaluation
2	Challenge 3. Cross-layer BPaaS Monitoring & Adaptation
2	Challenge 4. Intelligent Allocation of BPaaS across cloud levels

Area 8: Service Discovery and Composition

The discovery and composition of cloud services to satisfy customer requirements is still a complex and tricky task, requiring care and skill owing to the huge number of Cloud services which are currently available on the market.

Most of the Cloud Applications are composite and orchestrated solution realized through the integration of services. The orchestration of such composed solution has to enable to configure and coordinate the interaction among services automatically in Cloud environments. The orchestration is never an easy task, even more in cloud computing environments where it involves interconnecting processes running across heterogeneous systems in multiple locations, usually with proprietary interfaces. In particular, realizing sophisticated cloud application requires a cloud control framework that can orchestrate cloud resource.

Recently the concept of Cloud Pattern merged as a way to describe the composition and orchestration of Cloud Services in order to satisfy particular application requirements.

Such an automatic discovery and composition could be based on well-defined patterns in order to ensure the efficiency of the solution as well as on semantics in order to ensure the accuracy of the discovery activity and thus the robustness and suitability of the solution. Both activities should rely on both the functional and non-functional aspects as well as be able to scale to cater for the discovery and composition of millions of cloud services.

Associated Future Research Challenges

Challenge 1. Automatic discovery and composition of services

Automatically discover and compose cloud service at different levels (e.g., business process, software, infrastructure) in order to satisfy application or business requirements which will enable not only the fast development of applications and business processes for the cloud but also their runtime adaptation, when the respective need arises, thus catering for dynamicity. Such an automatic discovery and composition could be based on well-defined patterns in order to ensure the efficiency of the solution as well as on semantics in order to ensure the accuracy of the discovery activity and thus the robustness and suitability of the solution. Both activities should rely on both the functional and non-functional aspects as well as be able to scale to cater for the discovery and composition of millions of cloud services. Scalability could be achieved by taking a decentralized approach for cloud service discovery and composition that focuses on the distribution of information and computational load. Cloud service composition should be able to function even in the case of incomplete information and knowledge. Context-awareness is another aspect that has to be catered in order to enable a more effective composition of cloud services that better satisfies the requirements posed for the application or business process at hand.

Challenge 2. Automatic API Alignment and Software-defined everything

The automatic alignment among different APIs aims to discover and establish correspondences among the APIs elements that offer the same functionality. This feature will enable automatic porting of application among different cloud platforms without manual code rewriting. The combination of these advances together with innovative capabilities at multi-cloud management level in which the diversity of underlying resources are virtualized and therefore can be combined at programmatically level, will require of new means to provide advanced flexibility and manageability to multi-cloud set-ups.

Prioritisation

	Challenge
5	Challenge 1. Automatic discovery and composition of services
4	Challenge 2. Automatic API Alignment and Software-defined everything

Area 9: Dynamic Configuration, Provisioning, and Orchestration of Cloud Resources

Further evolution of multi-cloud models linked to cloud interoperability developments, leads to a richer Cloud ecosystem in which Cloud Broker manages service negotiations and relationships among Cloud consumers and providers, acting as an intermediary. It deals with request, performance tuning and delivery of services. The broker aggregates services from multiple providers, selecting the delivering agency according to some previously assigned score. However, advances are expected to go beyond the accepted capability of aggregating and brokering in a wide ecosystem of available commodity services that can be combined and assembled to deliver added-value. In this context, Broker serves as a service aggregator by integrating multiple existing services in order to deliver new services or functionalities. Data integration and security during transfers across multiple providers need to be assured by the broker itself.

Multi-cloud application deployment engines, allow definition of complex applications and their automated deployment onto multiple clouds. The adoption and extension of this type of tools will allow the selection and provisioning of resources based on user-defined algorithms. Moreover, these types of engines can be enabled with pluggable monitoring services to be deployed with an application to orchestrate the topology of the application dynamically while it is running. However, the advent of innovative cloud typologies, in form of fog computing set-ups or local clouds will make that today’s concept of Hybrid Clouds, goes beyond models in which a Cloud broker it is a central actor that manages provisioning and orchestration of resources across diverse Cloud set-ups, but complementarily move to decentralised and autonomic approaches in which application workloads define self-organisation, self-management and self-healing across a diversity of cloud deployments.

Associated Future Research Challenges

Challenge 1. Cross-Layer Cloud Service Negotiation

When requiring to realise the functionality of a cloud-based application or business process, there is a need to perform negotiation at different layers, especially when a respective broker is involved: (a) one layer considers the business/application requirements and involves the cloud broker and client as the negotiation parties and (b) multiple layers are involved in the negotiation for the cloud services (i.e., SaaS and PaaS) that constitute the user application or business process which will involve the cloud broker and the cloud provider of these services; (c) takes into account network awareness. In this respect, we complex negotiation has to be conducted which needs to be performed in a coordinated manner across the different layers and heterogeneous resources such that the satisfaction of the application/business requirements is guaranteed. Such complex negotiation should take into account the negotiation capabilities between the different participants in order to employ the most suitable negotiation protocols. It should also rely on transactional workflow techniques in order to support the coordination across the negotiations at different levels as well as guarantee that either all negotiations succeed or no agreement is reached in the end. The result of this complex negotiation should be either a composite SLA or a set of SLAs which depend on each other.

Challenge 2. Cloud Broker specialization for addressing specific vertical sector needs

Need for vertical market specific add to existing cloud service brokerage and market places capacities to consider specific vertical market needs in support enhancement of security, networking and specific regulations and compliance frameworks.

Challenge 3. Multi-Cloud improved application assembly and automation

Both in terms of generic cloud capabilities to deliver ubiquitous and consistent configuration and instrumentation across Cloud offerings, as well as sector-specific capabilities to tie together complex multi-tier and multi-tenant application delivery.

Challenge 4. Novel decentralized Inter-cloud computing continuum

Existing Inter-Cloud computing developments have to go beyond current centralised paradigm, to a pure Hybrid in a wider sense approach. This approach has to take into account diversity in cloud typologies and multi-cloud decentralised and autonomic management models, connecting a wider diversity and variety of entities and typologies of resources requiring new developments in a novel Inter-cloud computing continuum addressing the needs of specialised clouds for large data storage, analytics and edge local clouds for IoT.

Challenge 5. Self-* across a diversity of cloud deployments

The emergence of the above mentioned decentralised muti-cloud deployments will require that existing mechanisms for multi-cloud management evolve towards autonomic management systems that consider self-organisation, self-management and self-healing across a diversity of cloud deployments and coping with requirements of fault-tolerance substantiation, extreme scalability and parallelisation.

Challenge 6. Novel Orchestration and placement methods for hyper distributed cloud

Orchestration and placement problems in this hyper distributed cloud computing require the development of novel methodologies and technologies that take into account heterogeneity and trade-offs among consistency, reliability and performance.

Prioritisation

	Challenge
4	Challenge 3. Multi-Cloud improved application assembly and automation
4	Challenge 5. Self-* across a diversity of cloud deployments
4	Challenge 2. Cloud Broker specialization for addressing specific vertical sector needs
3	Research Area 9. Challenge 4. Novel decentralized Inter-cloud computing continuum
3	Challenge 6. Novel Orchestration and placement methods for hyper distributed cloud

Overview of Research Areas and challenges prioritisation

The cluster has done initial work in order to prioritise the identified Research Areas. The analysis has classified Research Areas according to Business Impact and Timeframe for realisation. In addition, priority of the Research Area as a whole has been assessed based on priority of the associated challenges. This process has been performed by a survey completed by cluster participants.

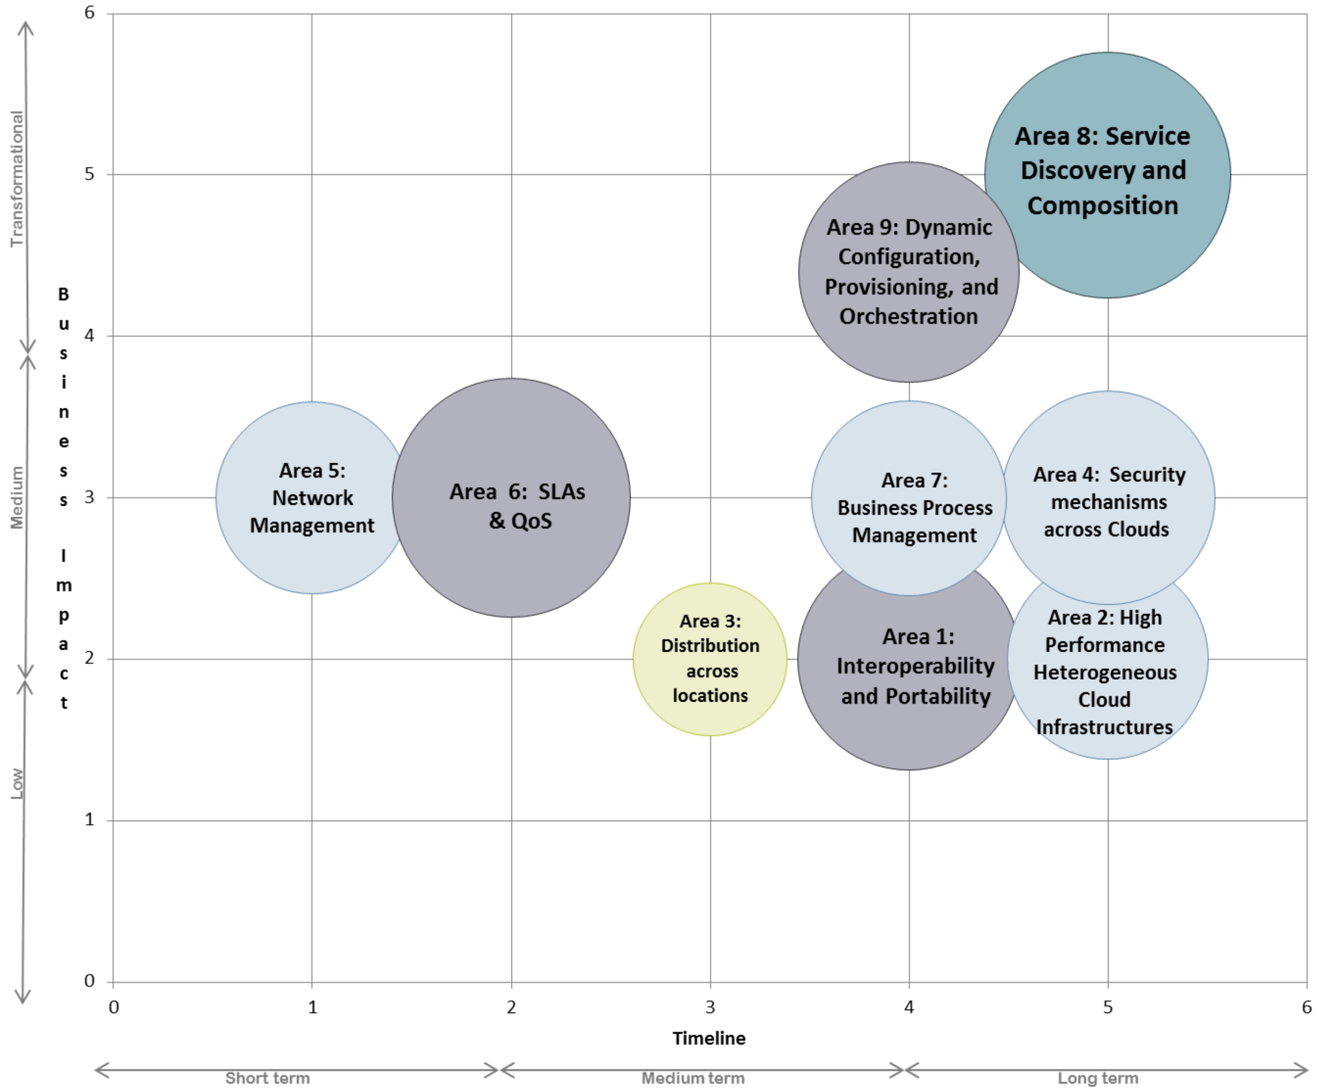

Findings of this analysis show that the area identified with major business impact in long term realisation is "Area 8 Service Discovery and Composition", considering the automatic discovery and composition of cloud services at different levels and taking into account scalability, decentralisation and automatization enabled by software defined everything. Although the expected business impact makes it feasible that market evolution alone will bring the realisation of Area 8 associated challenges, R&D investment would help European industry and research institutions to be well positioned for this expected market evolution.

This is followed, and in fact, intrinsically related to "Area 9 Dynamic Configuration, Provisioning and Orchestration of Cloud Resources" considering Hybrid Cloud developments and Cloud Brokerage realisations to develop in a wide ecosystem of cloud typologies and models following autonomic self-* approaches. Developments in both Area 8 and 9 are perceived as Transformational in order to flourish and evolve European Cloud market in the context of Digital Single Market for Europe, both providing new market opportunities for a diversity of cloud actors of all sizes including, among others, software vendors, cloud providers and telecom operators.

"Area 6 SLAs and QoS" is predicted to have shorter term market adoption, given current situation in which this is a major concern for wider cloud adoption. However, the expected evolution of SLAs in public Cloud providers will bring new challenges in the Inter-cloud field such as intelligent brokerage including composition of SLAs across diverse set of cloud providers and dynamicity on its management. Shorter term market adoption reflects the need of addressing the foreseen research needs in an earlier timeframe and the need of producing more mature research results that make use of several existing works in the area. While "Area 1 Interoperability and Portability" is logically one of the areas of concern in Inter-Cloud, which has the need of switching services among cloud typologies and providers transparently as its main challenge. Further research development of this area will have to rely and contribute to, not only on adoption of existing efforts in Cloud standardisation, but also in new developments in standardisation in Edge Computing.

Areas 5, Network management, 7, Business process management and 4, Security Mechanisms across clouds together with Area 2 High Performance Heterogeneous Cloud infrastructures are perceived as medium priority and medium business impact. . Nevertheless, it is expected that "Area 5 Network management" has a shorter adoption timeframe given maturity of related technologies such an NFV and SDN, even though additional research it is required to support additional fluidity in network management to edge and clouds levels.

Area 3 dealing with workload distribution across location for latency reduction, high availability and addressing legal constrains, it is perceived in this analysis as the area with lesser priority; this can be explained by related mechanisms that major public cloud providers are already articulating in the market and which are expected to be developed by themselves in the near future.

In addition to the analysis of priorities among Research Areas, it is important to remark that all proposed Research Areas have its roots in already developed and on-going research works. Considering these it is significant in order to allow future convergence of research results, research programmes' cohesion, as well as, overall resources optimisations. Taking this into account, and aiming to provide a clear link with the cluster's projects detailed in section 3 Analysis of EU funded research projects, the following table provides indications on the relation between proposed research areas and cluster's on going works:

Research Area	Previous Works in the area
Area 1: Interoperability and Portability	CloudSocket, ENTICE, ModaClouds, mOSAIC, PaaSage, SeaClouds, SSICLOPS, SWITCH
Area 2: High Performance Heterogeneous Cloud Infrastructures	ASCETiC, BEACON, CLOUDLIGHTNING, CloudSocket, CYCLONE, ENTICE, ModaClouds, mOSAIC, PaaSage, SeaClouds, SSICLOPS, SWITCH
Area 3: Distribution across locations	CloudSocket, BEACON, ENTICE, SeaClouds, OPTIMIS
Area 4: Security mechanisms across Clouds	CloudSocket, BEACON,OPTIMIS, ENTICE, CYCLONE
Area 5: Network Management	BEACON, CYCLONE, SSICLOPS, SWITCH
Area 6: SLAs & QoS	ASCETiC, mOSAIC, ModaClouds, SeaClouds, , SWITCH, OPTIMIS, CLOUDLIGHTNING
Area 7: Business Process Management	CloudSocket, mOSAIC, SWITCH
Area 8: Service Discovery and Composition	ENTICE, CYCLONE, mOSAIC, SWITCH
Area 9: Dynamic Configuration, Provisioning, and Orchestration	CLOUDLIGHTNING, ModaClouds, mOSAIC, PaaSage, OPTIMIS, SeaClouds

Although this document has structured Research challenges according to Inter-Cloud Research Areas identified, it is evident existing deep relations among them, and even, considerations of overlapping and dependencies. So, Research Areas and associated Challenges cannot be treated as separated silos, but instead recognizing their influences and relationships. For doing this, the following table provides an overview of prioritization of challenges across research Areas.

Priority	Research Area	Challenge
5	Area 1	Challenge 2.Switch services among cloud typologies and providers without efforts
	Area 2	Challenge 2.Monitor and guarantee inter-cloud infrastructure SLAs performance
	Area 4	Challenge 5.Auditability in Cloud Federated Cloud Networks
	Area 6	Challenge 3.Intelligent Broker
	Area 8	Challenge 1.Automatic discovery and composition of services
4	Area 1	Challenge 1.Develop once deploy everywhere
	Area 1	Challenge 3.Interoperability to cope with Cloud heterogeneity and application mobility
	Area 1	Challenge 6.Universal Semantic Service Description
	Area 2	Challenge 4.New languages to express overall high performance including storage, compute, network
	Area 4	Challenge 1.Security mechanisms for application integrity
	Area 4	Challenge 4.Definition of Security and network-aware application requirements
	Area 5	Challenge 2.To guarantee new paths for optimizing transfer of data among clouds, among IoTs and clouds-IoTs
	Area 6	Challenge 1.SLA Standard Representation
	Area 6	Challenge 2.Monitoring of QoS and application level monitoring
	Area 6	Challenge 4.SLA-based cloud service/application management
	Area 7	Challenge 1.Smart business-to-IT alignment
	Area 7	Challenge 2.Cross-layer and Scalable Multi-Cloud Workflows and BPaaS
	Area 8	Challenge 2.Automatic API Alignment and Software-defined everything
	Area 9	Challenge 2.Cloud Broker specialization for addressing specific vertical sector needs
	Area 9	Challenge 3.Multi-Cloud improved application assembly and automation
	Area 9	Challenge 5.Self-* across a diversity of cloud deployments
3	Area 1	Challenge 4.Automatic migration of in house application to the Cloud and across cloud typologies
	Area 2	Challenge 5.research area on Trust and privacy
	Area 5	Challenge 1.Extension to Cloud Federation concept and tools
	Area 5	Challenge 3.Enablement of responding more fluidly to changes in user demand at inter-cloud level but also at the edge level
	Area 7	Challenge 6.Flexible Cost Models
	Area 9	Challenge 4.Novel decentralized Inter-cloud computing continuum
	Area 9	Challenge 6.Novel Orchestration and placement methods for hyper distributed cloud
2	Area 1	Challenge 5.Extended Workload Portability
	Area 2	Challenge 3.Dynamic workload balancing in multi-cloud context
	Area 3	Challenge 1.Scalability across clouds based on demand
	Area 3	Challenge 3.Novel High Availability mechanism across hybrid cloud models
	Area 3	Challenge 4.Legal aspects
	Area 4	Challenge 2.Federated Authentication for non-Browser HTTP Applications
	Area 4	Challenge 3.Federated Authorization Policies and Use Cases
	Area 7	Challenge 3.Cross-layer BPaaS Monitoring & Adaptation
	Area 7	Challenge 4.Intelligent Allocation of BPaaS across cloud levels
	Area 7	Challenge 5.Smart Business Intelligence through cross-layer BPaaS Evaluation
1	Area 2	Challenge 1.Enable with inter-Cloud Service Provider connectivity
	Area 3	Challenge 2.Cross-cloud VM/container image distribution SLAs
	Area 5	Challenge 4.DevOps Agile development and deployment considering network management

3. Analysis of EU funded research projects

Before introducing in detail project's participating in Inter-cloud Challenges, Expectations and Issues Cluster, an overview of approaches followed for interoperability, technical topics, scenarios, standards and use cases is provided, so to present terms in classifications provided.

Classification

Inter-operability approach

Inter-operability approach defines mechanisms used to enable workloads movement among diverse cloud technologies and offerings. Projects participating in this cluster can be classified according in the following three criteria, as defined in [10].

- Cross-platform APIs: These provide a unified interface that abstracts diverse cloud management technologies and provider's APIs features and services. Examples of widely used cloud cross platform APIs are: jclouds [11], Deltacloud[12] and Libcloud[13].
- Semantics: Semantic Web technologies enable to define in a machine readable format common expression of services, APIs, resources, etc. across cloud technologies and providers. This mechanism commonly based in modelling of Cloud concepts based on OWL.
- Model-Driven: Model-driven software development methodologies allow developers to abstract from the specific cloud technology or platform of choice enabling decoupling application development and its deployment and operation.

Table 1 details the interoperability approaches that are considered in the projects participating in the cluster.

Inter-operability approach	Cross platform APIs	Semantics	Model-driven
CloudSocket	✓	✓	✓
BEACON	✓		
ENTICE		✓	✓
CYCLONE	✓		
ASCETiC	✓		✓
ModaClouds	✓		✓
SeaClouds	✓		
mOSAIC	✓	✓	
SWITCH	✓	✓	✓

Topics addressed

This subsection aims to identify specific topics addressed by projects in this cluster. Classification in topics addressed is makes use of Inter-cloud taxonomy defined in [5]. This classification considers the following categories for positioning cluster project's topics:

Provisioning mechanisms considered: Discovery, defines automated detection of services from a cloud provider; Selection, details automated choice of target cloud provider based on any defined criteria; Allocation, refers to resource planning capabilities at target cloud provider.

Topics addressed	Provisioning	Discovery	Selection	Allocation	Other
CloudSocket		✓	✓	✓	
BEACON			✓	✓	
ENTICE		✓	✓	✓	
CYCLONE		✓	✓	✓	Applications Deployment
ASCETiC			✓	✓	
ModaClouds			✓	✓	
SeaClouds		✓	✓	✓ (resource allocation at deployment time)	
mOSAIC		✓	✓	✓	
SWITCH		✓	✓	✓	

Level of workload and data portability, it considers Portability across providers; VM Portability, VM Mobility in a cloud platform, Data Portability and Application portability.

Topics addressed	Portability	Portability	VM Portability	VM Mobility	Data Portability	Application portability
CloudSocket			✓		✓	✓
BEACON						
ENTICE		✓	✓	✓	✓	✓
CYCLONE						
ASCETiC			✓	✓		
ModaClouds		✓			✓	✓
SeaClouds		✓	✓		✓ (partially)	✓
mOSAIC		✓				✓
SWITCH		✓	✓	✓	✓	✓

Service management capabilities offered, considering capacity to manage services across cloud providers and to effectively compare cloud service offerings.

Security features offered by project including Authorization and Identity management across cloud providers, whether Trust assurance mechanisms are part of the project, as well as, Federated Policy management mechanisms in aggregated services environment are considered.

Monitoring services of applications, as well as, physical and virtual resources.

Topics addressed		Multi-Cloud Service Management	Service Comparibility		Autorization and Identity	Trust	Federated Policy	Other		Monitoring	Other
CloudSocket	Service Management	✓	✓	Security	✓				Monitoring	✓	
BEACON		✓			✓						
ENTICE					✓	✓				✓	
CYCLONE		✓			✓	✓	✓	Brokering		✓	
ASCETiC		✓								✓	Energy parameters
ModaClouds		✓	✓							✓	
SeaClouds		✓	✓							✓	
mOSAIC		✓	✓		✓					✓	
SWITCH		✓	✓		Not directly	Not directly	Not directly			✓	

Network capabilities supported, including Connectivity, Addressing, Naming and Multi-casting issues.

Topics addressed		Connectivity	Addressing	Naming	Multi-casting	Other
CloudSocket	Network					
BEACON		✓				network federation, software-defined networking, network function virtualisation
ENTICE						
CYCLONE		✓				Network Abstraction Network management Network Isolation Resources Multitenancy Network Service Modelling
ASCETiC						
ModaClouds						
SeaClouds						
mOSAIC						
SWITCH		✓	✓	✓	✓	Software defined networking Time critical communication

Support for autonomic behaviour: This is, if the project considers self-managing capabilities based on i.e. any defined criteria or optimization problem.

SLA management features: In this context it identifies capabilities such as the aggregation of SLAs among different providers, multi-provider SLA; capability to federate these as well as, the ability to monitor and enact these SLAs across diverse providers.

Topics addressed		Autonomics		Multi-provider SLA	Federated SLA	SLA Monitoring and Dependency
CloudSocket	Autonomics	✓	SLA			✓
BEACON						
ENTICE					✓	✓
CYCLONE						
ASCETiC				✓		✓
ModaClouds		✓				✓
SeaClouds						✓
mOSAIC		✓		✓	✓	✓
SWITCH		✓		✓	✓	✓

Economic aspects considered: Market refers to trading features across providers, as well as pricing, accounting and billing features.

Business Process Management: It refers to the ability to spread and co-ordinately execute business processes across cloud multiple providers.

Topics addressed		Market	Pricing	Accounting and billing		Business Process Management
CloudSocket	Economy		✓	✓	Business Process Management	✓
BEACON						
ENTICE		✓				
CYCLONE						
ASCETiC			✓	✓		
ModaClouds						
SeaClouds		✓ (explore Cloud providers business models)	✓			
mOSAIC						✓
SWITCH		✓	✓			✓

Energy Efficiency: Consideration of Energy efficiency as a criteria to select providers and to optimize application execution.

Geo-location: Consideration of diverse geographical locations and if applicable, legal constrains.

Topics addressed		Geo-location	Legal issues		Energy efficiency
CloudSocket	Geo-location	✓		Energy Efficiency	
BEACON		✓			
ENTICE		✓	✓		✓
CYCLONE					
ASCETiC					✓
ModaClouds					
SeaClouds		✓			
mOSAIC					
SWITCH		✓			

Scenarios

As introduced in section 2, there are a diversity of multi and inter-cloud scenarios. This section aims to identify brokering and federation scenarios supported by project's participating in the cluster depending on their nature: provider centric or user centric. But also considering triggers for these scenarios: SLA based, scalability based and self-actions, considering autonomic behaviour.

Scenarios		Discovery	Selection	Allocation		Portability	VM Portability	VM Mobility	Data Portability		SLA	Scalability	Self-* actions
CloudSocket	Provider Centric	✓	✓	✓	User Centric		✓		✓	Trigger	✓	✓	
BEACON			✓	✓									
ENTICE		✓	✓			✓	✓	✓	✓		✓	✓	✓
CYCLONE		✓	✓	✓									
ASCETiC		✓	✓	✓		✓	✓	✓			✓		✓
ModaClouds			✓	✓		✓			✓		✓	✓	✓
SeaClouds		✓	✓	✓		✓	✓	✓	✓ (Partially)		✓	✓	✓ (automatic scaling; follow the sun)
mOSAIC		✓	✓	✓		✓					✓	✓	✓
SWITCH		✓	✓			✓	✓	✓			✓	✓	✓

Standards

This section identifies Cloud standards used in the different projects that take part in this project's cluster.

Standards	OVF	TOSCA	CIMI	CDMI	OCCI	CAMP	IEEE Standard for Intercloud Interoperability and Federation (SIIF)	Intercloud Architecture for Interoperability and integration	Other
CloudSocket									BPMN OWL RDF DMN CAMEL
BEACON			✓	✓	✓		✓		
ENTICE	✓			✓	✓				
CYCLONE		✓	✓						
ASCETiC	✓	✓							
ModaClouds		✓							
SeaClouds		✓				✓			
mOSAIC		✓			✓	✓			
SWITCH									Standards to apply under investigation, not decided yet

Use Cases

Due to the high number of variables that come into play in a complex cloud computing solution that involves portability and interoperability capabilities, a number of use case scenarios have been defined in the literature to underline requirements and consideration of the particular case.

The use cases labeled with the acronym “CSCC” are presented in the draft document *Interoperability and Portability for Cloud Computing: A Guide [14]* produced by the Cloud Standards Customer Council. The use cases branded with the acronym CCUS have been defined by the Cloud Computing Use Case Discussion Group [15].

CSCC S1: Customer Switches Providers for a Cloud Service: This scenario addresses the case of a customer currently using a Cloud Service provided by provider A, who wishes to switch to an equivalent service from provider B. This scenario touches many of the issues associated with portability.

From the point of view of Application code, the portability strategies depend on the Cloud service level: in the case of SaaS, the application code typically belongs to and is managed by the provider; in the case of PaaS, portability depends on the programming languages supported by the involved platforms and on the APIs offered by the cloud platform services to manage applications (submit the app code, configure, run ad control the app).

Data portability aspects must be considered too. At IaaS level, storage functionalities provided by Cloud vendors are typically low level, such as providing volumes for binary files or object storages, so customers are generally free to use their preferred data format.

While the same sometimes occurs at PaaS level, the situation tends to be more complex: a PaaS service provider may offer ready-to-go instances of databases, which may be sensitive to the data format chosen by the customer. However, there are some very generalized formats (CSV, XML, etc) which are supported by many types of database.

For SaaS services, data formats and contents are handled by the service provider, so major data portability considerations are needed.

CSCC 2: Customer Uses Cloud Services from Multiple Providers: This scenario addresses the case of a cloud service customer using cloud service 1 from Provider A and cloud service 2 from Provider B, while needing to use them together to achieve business goals.

Despite the benefits of mitigating the risk of data loss and temporary service unavailability, working with multiple cloud service providers can introduce logistical problems. Open source or vendor-agnostic tools and cloud management services can help solve these issues.

Not all cloud providers offer the same services. Storage and virtual machine (VM) instances are often standard, but services such as messaging or workflow and administration tools may vary across vendors. Working with multiple cloud providers may force customers into working with the lowest common denominator of services.

Sidestepping this problem is possible by choosing vendor-agnostic software applications that will run in all of the target clouds.

CSCC 3: Customer Links One Cloud Service to another Cloud Service: This scenario addresses the case of an architecture linking cloud services together to support a single application or an integrated set of applications. The advantage of multi-cloud deployments are several, first of all the ability of enterprises to leverage cloud solutions best fitting their stated needs, remaining within imposed cost limits.

An example of such a linking is represented by the case of a SaaS application capable of delivering basic business-related functionalities as needed by a customer, but that is not able to provide advanced functionalities (i.e. data analytics or business intelligence related). The customer can leverage IaaS capabilities from another cloud service provider, by migrate the data from the SaaS to the IaaS solutions, also in order to combine different data sets and perform more advanced analytics.

CSCC 4: Customer Links In-house Capabilities with Cloud Services: As more enterprises are planning their cloud investment, they will realize how they will leverage their existing in-house IT with their future cloud setting. A proper analysis of the available APIs of the both the in-house and cloud service will be required to understand how the integrated system will function and perform during typical execution.

CSCC 5: Migration of Customer Capabilities into Cloud Services: This scenario addresses the case of a customer, currently running an application or service on-premises, who wants to move that capability to a public cloud environment.

For SaaS cloud services migrating an on-premises application or service to a public cloud service provider does not involve porting the application code, because the application is being replaced. What is important in the SaaS case is the compatibility of the functional interface of the application, of any interfaces presented to end users and also any APIs made available to other customer applications. In order to reduce undesired side effects, the APIs made available by the SaaS service should be interoperable with the interface provided by the on-premises application or service that is being replaced. If the APIs are not interoperable any customer applications using the APIs will need to be changed as part of the migration process.

To migrate on a PaaS cloud, the application must be designed for one of the runtime environments available in the target PaaS service.

Generally a PaaS solution provides the elements of the particular software stack required by applications such as the operating system, an application server and a database, so that the customer only has to be concerned with the specific application components and data. Some concerns may arise regarding particular configurations required by the application, such as the ability to run scripts and the presence of certain tools for set-up, reporting or monitoring.

At IaaS level, the entire software stack is migrated through one or more virtual machine (VM) images, which can then be copied into the cloud service and executed there.

Some concerns arise if the application makes use of specialized device drivers or hardware devices that are unlikely to be supported by an IaaS provider.

CCUC 1: Changing SaaS vendors: This scenario addresses the case in which a cloud customer changes SaaS vendors. The data handled by one vendor's software should be importable by the second vendor's software, that means that both applications need to support common formats. Standard APIs for different application types will also be required.

CCUC 2: Changing middleware vendors: This scenario addresses the case in which a cloud customer changes cloud middleware vendors. Existing data, queries, message queues and applications must be exportable from one vendor and importable by the other. The requirement to achieve this porting is a common API for Cloud Middleware.

Cloud database vendors have enforced certain restrictions to make their products more elastic and to limit the possibility of queries against large data sets taking significant resources to process. For example, some cloud databases don't allow joins across tables, and some don't support a true database schema. Those restrictions are a major challenge to moving between cloud database vendors, especially for applications built on a true relational model.

CCUC 3: Changing VM hosts: This simple scenario addresses the case in which a cloud customer wants to take virtual machines, built on one cloud vendor's system, and run them on another cloud vendor's system. The main requirement of this operation is a common format for virtual machines.

CSC-CB: Cloud Bursting: The Cloud Bursting use case, focusing on Interoperability issues at IaaS level, describes a scenario where multiple Cloud Platforms need to work together. In particular, being it the typical situation, the use case describes the collaborations between public and private clouds and the possibility to move work and data loads between them.

The use case illustrates a situation in which a private cloud, running one or more Virtual Machines, needs more computational power from a public cloud in order to respond to a peak of incoming request from customers or to speed-up computation. In order to do so, VMs are dynamically migrated from one environment to another.

The use case is similar to the *CCUC 3: Changing VM hosts* scenario, but it differs from it not only for the dynamic requirements imposed, but also because the migrated VMs are supposed to collaborate with the one still hosted on the source Cloud platform.

			Cloud Standards Customer Council					Cloud Computing Use Case Discussion Group			
			CSCC S1: Customer Switches Providers for a Cloud Service	CSCC 2: Customer Uses Cloud Services from Multiple Providers	CSCC 3: Customer Links One Cloud Service to Another Cloud Service	CSCC 4: Customer Links In-house Capabilities with Cloud Services	CSCC 5: Migration of Customer Capabilities into Cloud Services	CCUC 1: Changing SaaS Vendors	CCUC 2: Changing Middleware Vendors	CCUC 3: Changing VM Hosts	CSC-CB: Cloud Bursting
Interoperability	IAAS	Data	CYCLONE MODAClouds SWITCH	ENTICE CYCLONE MODAClouds SWITCH	CYCLONE SWITCH	ENTICE CYCLONE MODAClouds SWITCH	CYCLONE SWITCH	CYCLONE	CYCLONE	ENTICE CYCLONE SWITCH	ENTICE CYCLONE SWITCH
		Service	CYCLONE ASCETiC MODAClouds SeaClouds mOSAIC SWITCH	ENTICE CYCLONE MODAClouds mOSAIC SWITCH	ENTICE CYCLONE mOSAIC SWITCH	CYCLONE ASCETiC MODAClouds SeaClouds mOSAIC SWITCH	CYCLONE mOSAIC SWITCH	CYCLONE SWITCH	CYCLONE SWITCH	ENTICE CYCLONE ASCETiC SWITCH	ENTICE CYCLONE ASCETiC mOSAIC SWITCH
		Application	CYCLONE ASCETiC MODAClouds SeaClouds mOSAIC SWITCH	BEACON (BEACON applications can be deployed in a cloud federation.) ENTICE CYCLONE MODAClouds mOSAIC SWITCH	ENTICE CYCLONE mOSAIC SWITCH	CYCLONE MODAClouds SeaClouds mOSAIC SWITCH	CYCLONE ASCETiC mOSAIC SWITCH	CYCLONE SWITCH	CYCLONE SWITCH	ENTICE CYCLONE ASCETiC SWITCH	ENTICE CYCLONE ASCETiC mOSAIC SWITCH
		System	CYCLONE ASCETiC MODAClouds mOSAIC SWITCH	BEACON (provides tools to manage IaaS cloud federations) ENTICE CYCLONE MODAClouds mOSAIC SWITCH	CYCLONE mOSAIC SWITCH	CYCLONE MODAClouds mOSAIC SWITCH	CYCLONE ASCETiC mOSAIC SWITCH	CYCLONE SWITCH	CYCLONE SWITCH	ENTICE CYCLONE ASCETiC SWITCH	BEACON /supports applicatyion deployment accross the different clouds of a cloud federation) ENTICE CYCLONE ASCETiC mOSAIC SWITCH
	PAAS	Data	CYCLONE SeaClouds (parcially) SWITCH	ENTICE CYCLONE SeaClouds (parcially) SWITCH	CYCLONE SWITCH	CYCLONE SeaClouds (parcially) SWITCH	CYCLONE SWITCH	CYCLONE	CYCLONE	CYCLONE SWITCH	CYCLONE SWITCH
		Service	CYCLONE SeaClouds mOSAIC SWITCH	ENTICE CYCLONE SeaClouds mOSAIC SWITCH	ENTICE CYCLONE mOSAIC SWITCH	CYCLONE SeaClouds mOSAIC SWITCH	CYCLONE mOSAIC SWITCH	CYCLONE	CYCLONE	CYCLONE SWITCH	CYCLONE SWITCH
		Application	CYCLONE SeaClouds mOSAIC SWITCH	ENTICE CYCLONE SeaClouds mOSAIC SWITCH	ENTICE CYCLONE mOSAIC SWITCH	CYCLONE SeaClouds mOSAIC SWITCH	CYCLONE mOSAIC SWITCH	CYCLONE	CYCLONE	CYCLONE SWITCH	CYCLONE SWITCH
		System	CYCLONE mOSAIC SWITCH	CYCLONE mOSAIC SWITCH	CYCLONE mOSAIC SWITCH	CYCLONE mOSAIC SWITCH	CYCLONE mOSAIC SWITCH	CYCLONE	CYCLONE	CYCLONE SWITCH	CYCLONE SWITCH
	SAAS	Data	CYCLONE	CYCLONE	CloudSocket CYCLONE	CYCLONE	CYCLONE	CloudSocket CYCLONE	CYCLONE	CYCLONE	CloudSocket CYCLONE
		Service	CYCLONE mOSAIC	CYCLONE	CloudSocket CYCLONE mOSAIC	CYCLONE mOSAIC	CYCLONE mOSAIC	CYCLONE	CYCLONE	CYCLONE	CloudSocket CYCLONE
		Application	CYCLONE mOSAIC	CYCLONE	CYCLONE mOSAIC	CYCLONE mOSAIC	CYCLONE mOSAIC	CYCLONE	CYCLONE	CYCLONE	CYCLONE
		System	CYCLONE mOSAIC	CYCLONE	CYCLONE mOSAIC	CYCLONE mOSAIC	CYCLONE mOSAIC	CYCLONE	CYCLONE	CYCLONE	CYCLONE

			Cloud Standards Customer Council					Cloud Computing Use Case Discussion Group			
			CSCC S1: Customer Switches Providers for a Cloud Service	CSCC 2: Customer Uses Cloud Services from Multiple Providers	CSCC 3: Customer Links One Cloud Service to Another Cloud Service	CSCC 4: Customer Links In-house Capabilities with Cloud Services	CSCC 5: Migration of Customer Capabilities into Cloud Services	CCUC 1: Changing SaaS Vendors	CCUC 2: Changing Middleware Vendors	CCUC 3: Changing VM Hosts	CSC-CB: Cloud Bursting
Portability	IAAS	Data	CloudSocket SeaClouds SWITCH	SeaClouds SWITCH	SWITCH	SeaClouds SWITCH	SeaClouds SWITCH			CloudSocket SWITCH	CloudSocket SWITCH
		Service	ASCETiC SeaClouds mOSAIC SWITCH	SeaClouds mOSAIC SWITCH	SWITCH	SeaClouds SWITCH	SeaClouds mOSAIC SWITCH			mOSAIC SWITCH	mOSAIC SWITCH
		Application	ASCETiC SeaClouds mOSAIC SWITCH	SeaClouds mOSAIC SWITCH	SWITCH	SeaClouds SWITCH	SeaClouds mOSAIC SWITCH			mOSAIC SWITCH	mOSAIC SWITCH
		System	ASCETiC SeaClouds mOSAIC SWITCH	SeaClouds mOSAIC SWITCH	SWITCH	SeaClouds SWITCH	SeaClouds mOSAIC SWITCH			mOSAIC SWITCH	mOSAIC SWITCH
	PAAS	Data	SeaClouds(parcially) SWITCH	SeaClouds(parcially)	SWITCH	SeaClouds(parcially) SWITCH	SeaClouds(parcially) SWITCH			SWITCH	SWITCH
		Service	SeaClouds mOSAIC SWITCH	SeaClouds mOSAIC	SWITCH	ASCETiC SeaClouds SWITCH	SeaClouds mOSAIC SWITCH		mOSAIC	SWITCH	SWITCH
		Application	SeaClouds mOSAIC SWITCH	SeaClouds mOSAIC	SWITCH	SeaClouds SWITCH	SeaClouds mOSAIC SWITCH		mOSAIC	SWITCH	SWITCH
		System	mOSAIC SWITCH	SWITCH	SWITCH	SWITCH	mOSAIC SWITCH		mOSAIC	SWITCH	SWITCH
	SAAS	Data	CloudSocket SWITCH	SWITCH	SWITCH	SWITCH	SWITCH	CloudSocket		SWITCH	CloudSocket SWITCH
		Service	mOSAIC SWITCH	SWITCH	SWITCH	SWITCH	mOSAIC SWITCH	mOSAIC		SWITCH	SWITCH
		Application	mOSAIC SWITCH	SWITCH	SWITCH	SWITCH	mOSAIC SWITCH	mOSAIC		SWITCH	SWITCH
		System	mOSAIC SWITCH	SWITCH	SWITCH	SWITCH	mOSAIC SWITCH	mOSAIC		SWITCH	SWITCH

EU funded research projects

CloudSocket Project

Overview

Although large enterprises benefit from cloud technology, SMEs are hampered due to missing IT-competence and hence lose the ability to efficiently adapt their IT to their business needs. The H2020 project CloudSocket investigates the idea to address customers' business needs with cloud offerings in form of Business Processes as a Service (BPaaS). Concept models and semantics are used to align domain specific business processes with executable workflows that are deployed and brought into production in a multi cloud environment. The Business Process Management System paradigm (BPMS) leads to requiring the realization of the functional capabilities of so-called BPaaS Environments, focusing on the four main phases of the business process lifecycle involving the: (i) design, (ii) allocation, (iii) execution and (iv) evaluation of a business process, which technically compose the CloudSocket Broker platform. Multi-cloud aspects are introduced in the execution environment via the need to place different BPaaS components in the cloud, such as a workflow engine and software components which realize the functionality of BPaaS tasks, to appropriately manage the respective data involved and the corresponding load and to adapt the deployment and selection of services to guarantee that the promised service levels are delivered to the BPaaS customers.

Approach to Inter-cloud

Inter-cloud aspects are approached by two main BPaaS Environments, namely the BPaaS Allocation Environment and the BPaaS Execution Environment.

The goal of the BPaaS Allocation Environment is to configure allocation directives and rules for an executable workflow model to be deployed and executed in the cloud. An executable workflow model, as produced by the BPaaS Design Environment, does not contain information about which concrete services can be exploited in order to realize the functionality of the business process tasks. As such, driven also by the business and technical requirements, the BPaaS Allocation Environment supports the CloudSocket Broker in making an informed selection of which services from the candidate ones to select for each business process task. The same set of requirements should also drive the decision about which IaaS offerings to select in order to deploy the BPaaS software components (e.g., internal services). Through both types of selection, the ending result would be not only a fully executable workflow model but also a deployment plan which will enable the deployment of the BPaaS, thus enabling its execution by the BPaaS Execution Environment. Another connecting piece related to a

BPaaS and its deployment refers to the specification of adaptation rules that can drive the adaptation of a BPaaS when such a need arises. Such rules are important if a more or less constant service level is exhibited by the BPaaS to its customers. The three main products of the BPaaS Allocation Environment are encapsulated in a so called BPaaS bundle which can then be published in the Marketplace in order to be available to the CloudSocket Customers. Thus, what can actually be externally exploited by any other environment is the models encapsulated by the BPaaS bundle.

Concerning the inter- & multi-cloud aspect, we can actually see that a BPaaS can be realized through the selection of cloud services at different levels (SaaS & IaaS) which can map to different cloud providers. This can be justified also by the fact that different business or technical requirements can lead to selecting not only one but many cloud providers, especially in cases where special features are needed which cannot be supported by just few cloud providers. Apart from the design-time selection of services possibly situated in different clouds, there is also the dynamic selection and adaptation aspect. This aspect is addressed at the information level by the Allocation Environment through the incorporation in the BPaaS bundle of prioritised rules, which dictate the way services can be selected for a particular BPaaS task according to the current context, as well as of adaptation rules which dictate what adaptation actions can be performed when particular problematic situations occur. As an example, we can have the case that the performance of a particular SaaS service is problematic and leads to the violation of an SLO. In that case, there can be an adaptation rule indicating that this service can be substituted by another one which is offered by a different cloud provider, when such performance degradation is detected.

The BPaaS Execution Environment aims at deploying and executing a BPaaS bundle, once this has been purchased by a customer at the BPaaS Marketplace. Thus, this environment actually takes care of deploying the BPaaS according to the deployment plan included in the bundle as well as importing the respective executable workflow model into a workflow engine in order to enable its execution by the customer that has purchased it. As such, the execution of the workflow encapsulated in the BPaaS bundle is supported. Another goal of this environment is to support the monitoring and evaluation of the BPaaS according to the KPIs and SLOs that have been defined for it. In case of KPI or SLO violations, particular adaptation plans are executed which are triggered via the adaptation rules that have been already defined in the BPaaS bundle Concerning the external functionality, as can be seen by a CloudSocket Customer and the other environments, the BPaaS Execution Environment exposes a functionality which enables the deployment of a BPaaS, the creation, execution and management of instances of the workflow encapsulated by the BPaaS and the production and support for retrieval of BPaaS monitoring and assessment results (for evaluation purposes in the BPaaS Evaluation Environment).

Concerning the inter- & multi-cloud aspect, we actually foresee two different but related capabilities of the BPaaS Execution Environment. The first capability is already mentioned above in the case of the BPaaS Allocation Environment. In particular, the BPaaS Execution Environment is the implementation instrument for dynamic selection and adaptation rules which can drive the re-deployment and re-configuration of the BPaaS in a multi-cloud setting, even if the initial setting is a single-cloud one. The second capability concerns the self-adaptation of the environment. There can be different cases where such a self-adaptation is needed and must be performed. Such self-adaptation can involve creating additional instances for the whole environment or some of its components or increasing the amount of resources for the already exploited VMs. Moreover, such self-adaptation will eventually make the distributed configuration and deployment of the BPaaS Execution Environment multi-cloud, especially for particular cases where e.g. different locations need to be covered for different clusters of BPaaS customers which cannot be serviced by the same cloud provider.

			Cloud Standards Customer Council					Cloud Computing Use Case Discussion Group			
Use Cases			CSCC S1	CSCC 2	CSCC 3	CSCC 4	CSCC 5:	CCUC 1	CCUC 2	CCUC 3	CSC-CB
Interoperability	IAAS	Data									
		Service									
		Application									
		System									
	PAAS	Data									
		Service									
		Application									
		System									
	SAAS	Data			X			X			X
		Service			X						X
		Application									
		System									
Portability	IAAS	Data	X							X	X
		Service									
		Application									
		System									
	PAAS	Data									
		Service									
		Application									
		System									
	SAAS	Data	X					X			X
		Service									
		Application									
		System									

The relevant use cases for CloudSocket are CSCC S1, CSCC S2, CSCC S3, CCUC1, CCUC2, and CSC-CB. CSCC S1 case can occur when the system decides to adapt a particular cloud service for any specific reason, such as the deterioration of service performance. Although this violates with the specification of the scenario as the switching is not customer-driven but system-driven, it is mentioned here as the respective issues involved are actually the same as those defined in the use case: (a) there are portability problems as data created for the former cloud provider, which are stored/represented in a particular format, must be probably transformed into the format required for the new cloud provider (SaaS case); (b) in case of switching between different database services (IaaS/PaaS), there is also the problem of data portability, especially when the databases involved are either of a different type (e.g., SQL vs NoSQL) or require some kind of special data manipulation or transformation (e.g., in case of data manipulation language heterogeneities). In addition, in order to be truly multi-cloud, there is also the basic need to be able to exploit the equivalent services offered by different cloud providers (especially at the IaaS level).

CSCC S3 is the actual main use case for CloudSocket as a BPaaS can be considered as a combination of different types of cloud services (IaaS & SaaS). In this sense, the main issues that are considered relevant and have to be addressed are service and data interoperability.

The major issues revolve around data and touch both data interoperability and portability. Data interoperability is required in cases where the output of one SaaS service S1 has to be used as input for a second SaaS service S2. When S2 is substituted by another service S3 offered by a different cloud provider, there can be the need to transform the output of S1 to the appropriate format which is expected for the input of S3. To deal with such a case, a special component called Data Mediator has been envisioned for CloudSocket which will be able to map as well as transform the data from service S1 to that of service S3 by exploiting well-known data mapping and transformation techniques. Data portability is required when the data exploited internally by the previous service have to be imported to the service substituting it. For instance, for email services, this can map to copying the email history and the address books of one or more users to guarantee that the same email experience is delivered when one email service is substituted by another one offered by a different cloud provider. Obviously, the exploitation of standards in the respective domains involved can play a crucial role as this will remedy the problem and just map it to the simple case of data export and import without any kind of transformation.

The change of VM hosts is certainly another adaptation-based scenario in CloudSocket. The consortium will need to address the issue of image portability which can be of course remedied via the adoption of common/standardized image formats.

Finally, the CSC-CB use case is relevant for CloudSocket as it can well map in cases where different cloud platforms have to cooperate in a coordinated fashion to accommodate for the additional load incurred to one or more BPaaS. This use case, in the context of CloudSocket, is related to issues of data and service interoperability (e.g., one service directly calls an external service, data has to be transformed to the format required for a service to be called next) as well as data portability (e.g., moving images) which have to be addressed in order to enable such coordinated cooperation between cloud platforms.

BEACON project

Overview

The BEACON H2020 project aims at enabling federated cloud networking. The recent development of software defined networking and network virtualization technologies has created the opportunity to fully integrate network virtualization technologies into cloud middleware. This will enable management of advanced hybrid clouds and heterogeneous cloud federations. Network virtualization technologies from the OpenDove project will be integrated with open source cloud middleware OpenNebula and OpenStack.

Approach to Inter-cloud

The BEACON project aims to enhance cloud middleware market with network virtualization technology to support the management of hybrid clouds and cloud federations. Our proposal will deliver a homogeneous virtualization layer, on top of heterogeneous underlying physical networks, computing and storage infrastructures, providing enablement for automated federation of applications across different clouds and datacentres. The figure below shows the BEACON federated cloud architecture. The service manager is responsible for the instantiation of the application by requesting the creation and configuration of virtual machines for each service component included in the service definition, using the interfaces exposed by the cloud manager. The Cloud Manager is responsible for the placement of VMs into VM Hosts. It receives requests from the Service Manager through the Cloud interface to create and resize VMs, and finds the best placement that satisfies a given set of constraints. The Cloud Manager is free to place, and move, the VMs anywhere, even on remote sites within the federation, as long as the placement satisfies the constraints. The network manager is responsible for allocating network resources to manage federated cloud virtual network and overlay networks across geographically dispersed sites. The right part of the figure shows a second cloud stack running on a different cloud provider. Together they form a federation with two cloud providers. The middle part of the figure shows that the cloud manager and network managers of the two cloud providers communicate to share resources and manage the federation.

The BEACON project aims to define a framework to federate networks across multiple heterogeneous clouds. To this end the project will integrate software defined networking (SDN) and network function virtualization (NFV) technology into cloud middleware so that virtualized compute and networking resources can be managed coherently. The network virtualization technology is OpenDove and the cloud middleware is OpenNebula and OpenStack. The BEACON framework will support federated cloud networking in both loosely coupled federations and tightly coupled federations. In particular the federation of networks will be supported in hybrid clouds where a network of a private cloud is federated with networks from one or more public clouds. The federated networks will be secured by security functions deployed as NFV across the cloud federation. The project will be validated with a flight scheduling service which is deployed in a hybrid cloud and serves customer airlines across the globe.

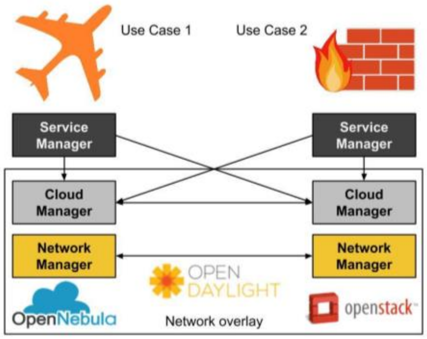

Fig. 1. BEACON Federated Cloud Networking Architecture

Cloud networking aspects will be based on OpenDove, a collaborative project under The Linux Foundation. We will extend the OpenDOVE project with new rich inter-cloud APIs to provision cross-site virtual network overlays. The new inter-cloud network capabilities will be leveraged by existing open source cloud platforms, OpenNebula and OpenStack, to deploy multi-cloud applications. Different aspects of the platforms will be extended to accommodate the federated cloud networking features like multi-tenancy, federated orchestration of networking, compute and storage management or the placement and elasticity of the multi-cloud applications.

Use Cases

			Cloud Standards Customer Council					**Cloud Computing Use Case Discussion Group**			
			CSCC S1	CSCC 2	CSCC 3	CSCC 4	CSCC 5:	CCUC 1	CCUC 2	CCUC 3	CSC-CB
Interoperability	**IAAS**	Data									
		Service									
		Application		x (BEACON applications can be deployed in a cloud federation.)							
		System		x (BEACON provides tools to manage IaaS cloud federations)							x (BEACON supports application deployment across the different clouds of a cloud federation)

SSICLOPS project

Overview

SSICLOPS focuses on techniques for the management of federated private cloud infrastructures, in particular cloud networking techniques, within software-defined datacentres and across Wide-Area Networks.

The project will design, implement, demonstrate, and evaluate four specific use cases; namely, cloud-based in-memory databases, the analysis of high-energy physics workloads, Network Function Virtualisation in a next generation Point of Presence, and caching and content delivery.

SSICLOPS will empower enterprises to create and operate high-performance private cloud infrastructures that allow flexible scaling through federation with other private clouds without compromising their service level and security requirements.

The SSICLOPS federation will support the efficient integration of clouds, no matter if they are geographically co-located or distributed, belong to the same or different administrative entities or jurisdictions: in all cases, SSICLOPS will deliver maximum performance for inter-cloud communication, enforce legal and security constraints, and minimize the overall resource consumption. In such a federation, individual enterprises will be able to dynamically scale in/out their private cloud services, because they dynamically offer own spare resources (when available) and take in resources from others when needed. This allows to maximize own infrastructure utilization while minimizing excess capacity needs for each federation member.

SSICLOPS-powered private clouds will offer fine-grained monitoring and tuning capabilities along with workload planning and optimization tools to maximize the performance across a broad spectrum of workloads and across a wide operational scale, as we will demonstrate using three highly diverse use cases. The SSICLOPS solution will be based upon state-of-the-art open source products used broadly in private cloud deployments today to provide enterprises with full control over their own deployment.

Approach to Inter-cloud

In order to enable the inter-cloud communications SSICLOPS developing secure network infrastructure and increasing the security levels for intra- and inter- cloud communication.

The inter-cloud computing is based on relationships (pattern) among multiple cloud service providers (CSP). This pattern (peering, federation or intermediary) allows the CSP to interworks with one or more peer CSPs to assure intermediation and secure of services provided by these CSPs.

The trusted inter-cloud relationship among multiple CSPs rely on confidence between cloud service customer (CSC) and CSP or between CSPs. One of them have to shift the physical control over application, service, resource and data to the others. The appropriate secure mechanisms (e.g. security access control, security of network connectivity between the CSPs) should be supported during peer CSPs interactions to achieve trusted inter-cloud computing rely.

The relevant CSPs are supposed to join a common trusted inter-cloud to establishment a trust relationship between them. In particular, the multiple CSPs involved in inter-cloud may be administrated by different parties. In case of an inter-cloud federation, the involved CSPs may establish trust relationships among them prior to any interactions between them or during inter-cloud interactions (e.g., service requests between CSPs). Therefore, the SSICLOPS is investigating solutions for:

- *Trusted Federated Cloud Computing* (specification and enforcement of security and privacy policies),
- *Security-Aware Storage and Processing* (enabling cloud data processing functions to become security-aware and apply appropriate security measure to sensitive data),

- *Hardened Cloud Data Transport* (securing network system platforms and control planes and lightweight/efficient security for (inter-cloud transport, and
- *Performance, reliability and flexibility* for intra- and inter-cloud transport.

The security in inter-cloud computing can be realize based on:

a) Self-Service Security which enable self-service management of security in heterogeneous cloud infrastructures and provide flexible mechanisms to let CSC or CSP control in a fine-grained manner the security of their resources in the cloud.

b) Self-Managed Security which enable full automation of security management in order to reduce operational costs while adding more flexibility and providing a unified view of security in heterogeneous cloud environments.

c) End-to-End Security which implement a distributed security abstraction layer between endpoints defined by CSC to overcome the heterogeneity of security technologies across clouds and to manage trust relationships between different layers and across CSPs, to provide a unified user experience of security.

The security and privacy of inter-cloud is the main challenge of integrating CSP platforms. This is necessary to provide self-service, self-managed and end-to-end security services for the CSC, and for the CSP to guarantee a level of confidentiality, integrity, and availability of resources hosted on CSPs clouds.

The security and privacy of inter-cloud is based on distributed cloud management. It enables the primary CSP to provide end-to-end dynamic deployment, configuration and unified control of security and resilience of cloud services across multiple CSPs. In implementation, distributed cloud management supported trust could be realized by combining specialized protocol design with smart interaction with the underlying cloud network fabric (e.g., using SDN traffic engineering and cloud-tailored smart queue management).

To increase security and privacy (including data integrity, localization and confidentiality) of inter-cloud computing, it is required to define a terminology (language) to annotate (or tag) workloads and data with security requirements (such as permissible storage locations). These annotations will be processed by the system during scheduling and migration to ensure that workload confines are kept. Additionally, annotation of workload allows the use of appropriate network (e.g. SDN) data plane mechanisms for strong security protection and traffic isolation to ensure that the above constraints are reached when workloads are practically placed, executed (data accessed and stored) and migrated.

Private and hybrid cloud technology is a key enabler for future ICT growth, especially when deployed in regions of the world with strong data confidentiality, privacy and consumer protection legislation such as the EU.

Hyperscaler public clouds have open APIs, but the system software of the cloud infrastructure is proprietary (and carefully – manually – optimized). On the other hand, private and hybrid clouds are usually put together based on open-source components (APIs are public and system software is public). Open source is good for security, transparency, and cost; however, cloud infrastructure assembled only from open-source systems and commodity hardware has a lower degree of integration and optimization compared to the custom code that runs the hyper scalers. That leads to a performance difference, and it can cause feature disparity, creating the undesirable incentive to migrate sensitive data from private to public clouds in search of improved reliability.

The goal of SSICLOPS is to minimize those performance and feature set differences by revisiting and re-optimizing the whole software infrastructure used in private clouds from applications, protocols, operating systems and networking stack and hypervisors. The result will be that private cloud infrastructures gain performance and feature parity with hyperscaler public clouds, a key challenge for inter cloud infrastructure technology.

ENTICE project

Overview

The ENTICE project aims to provide a ubiquitous repository-based technology called ENTICE environment, a universal backbone for IaaS VM management operations, which accommodate the needs for different use cases with dynamic resource (e.g. requiring resources for minutes or just for a few seconds) and other QoS requirements. This ENTICE technology is completely decoupled from the applications and their specific runtime environments, but continuously supports them through optimised VM image creation, assembly, migration and storage.

Approach to Inter-cloud

As depicted in Figure 1, the proposed management operations will be carried out in an inter-cloud environment composed of different cloud providers and operators.

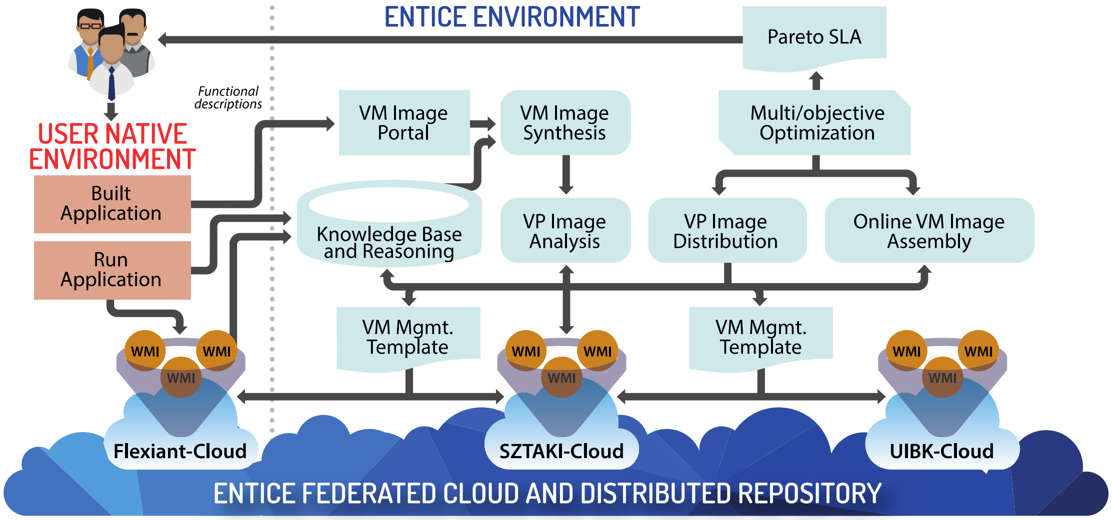

Figure 1: ENTICE overall architecture.

The main interests of ENTICE in the research field of cloud federation lie in its proposed federated VM image distribution strategies. Therefore data management (i.e., VMI handling and secure storage across clouds), issues are its primary concerns, and federation formation and operation is planned to be studied for data distribution. Thus, for ENTICE cross platform APIs and semantics are needed to handle VMI management in different clouds.

Although ENTICE does not tackle VM management related issues on its own it expects federated and inter cloud related basic provisioning functionalities for VM lifecycle management to operate some of its components (e.g. VM Image synthesis). Besides, VM image portability is and inherent property of the ENTICE environment as it aims at offering a generic technique to transform VMIs for the needs of the different clouds with several techniques (like classical image conversion and/or online VMI reconstruction). Different VMI formats and different hypervisors (virtualization techniques ranging from Xen to VMWare and even including docker) are planned to be supported by the ENTICE environment.

Pareto-SLAs are used to guarantee service deployment performance and elasticity for certain operations relevant for stakeholders (specially users – i.e. SaaS/PaaS providers on top of IaaS systems - and even for IaaS providers). Data management (image delivery) and VMI synthesis will be carried out by taking into account the geographical location of VMIs or their fragments (ENTICE aims at reducing storage costs and energy needs by fragmenting images, reducing their inherently replicated content and distributing them on a location aware way). For details, see ENTICE D2.3 Deliverable on ENTICE environment and architecture design.

Standards and quasi-standards related to VM image formats (e.g. OVF, AMI, docker) and image delivery among federation partners are taken into account for the design of the ENTICE environment and most of them are planned to be supported and used.

Use Cases

			Cloud Standards Customer Council					Cloud Computing Use Case Discussion Group			
			CSCC S1	CSCC 2	CSCC 3	CSCC 4	CSCC 5:	CCUC 1	CCUC 2	CCUC 3	CSC-CB
Interoperability	**IAAS**	Data		✓		✓				✓	✓
		Service		✓	✓					✓	✓
		Application		✓	✓					✓	✓
		System		✓						✓	✓
	PAAS	Data		✓							
		Service		✓	✓						
		Application		✓	✓						
		System									
	SAAS	Data									
		Service									
		Application									
		System									

The ENTICE environment will play an important role in the VMI related data management needs of cloud providers. It operates over IaaS systems and offers some PaaS capabilities, therefore these service levels are supported in the cloud application lifecycle. ENTICE does not handle interoperability issues related to VM management, only software deployment issues are addressed.

CYCLONE project

Overview

CYCLONE stands for: Complete Dynamic Multi-Cloud Application Management. Complex applications are often distributed between cloud infrastructures to provide resilience against cloud provider outages, higher levels of elasticity, and better response times for their users by placing services near the clients. Moreover, they are often designed to scale automatically in response to demand and to permit live upgrades of the underlying software.

The CYCLONE project primarily targets application service providers who develop complex computing platforms and deploy them on cloud infrastructures by bringing together Network-as-a-Service, application deployment, service access management and end-to-end security solutions for Multi-Cloud environments. One often finds these application service providers in innovative SMEs and institutes that want to take advantage of the cloud's dynamic nature to deliver responsive, adaptable services to their users.

The project partners have identified two flagship domains: academic use cases for bioinformatics research and use cases for a commercial deployment for smart grids in the energy sector that will guide the development of the tools in order to facilitate the deployment, management, and use of complex, multi-cloud applications, and enhance the end-to-end security of those applications. More specifically, CYCLONE will:

- **Improve cloud services in the Infrastructure-as-a-Service (IaaS)** layer by integrating network services into the cloud offering.
- Develop tools that provide enhanced functionality **for cloud providers** that agree to federate their resources.
- **Provide tools for application developers** to automate the placement of service components, scale resources toward a full-featured Platform-as-a-Service (PaaS) offering.
- Develop mechanisms to more easily deploy and manage applications and, thus, maintain **Software-as-a-Service (SaaS)** systems.
- **Provide software** that allows developers to ensure the end-to-end, secure use of data within their application as well as secured access to remote data sources.
- **Demonstrate** that the CYCLONE software meets the needs of concrete cloud-based services in federated, heterogeneous and multi-layered (IaaS, PaaS, SaaS) cloud environments. While providing frequent, production-quality releases of that software,

Approach to Inter-cloud

There are numerous software products that manage different aspects of cloud solutions, namely:

1. Application deployment and management software
2. Cloud security and access management software
3. Network management software

There is currently no existing Inter-Cloud solution to manage all of these areas. This lack of holistic management capabilities impacts especially multi-cloud deployments, as those are composed of a multitude of different instances of clouds, networks and services, which often have to be managed independently, thus creating considerable management overhead.

Thus, CYCLONE approach aims to integrate partner's established cloud solutions for managing software-defined networking, application deployment, cloud security and access management into a holistic cloud action and resource model. These integrated models create a holistic cloud management platform, which empowers platform users to deploy their services on any cloud of their choosing and still be able to manage it uniformly and bring added value to the Inter-Cloud. The concept is shown in Figure 2. Using multiple clouds can improve application reliability, application performance (e.g. if deployed near clients) and exploit different capabilities of different cloud platforms.

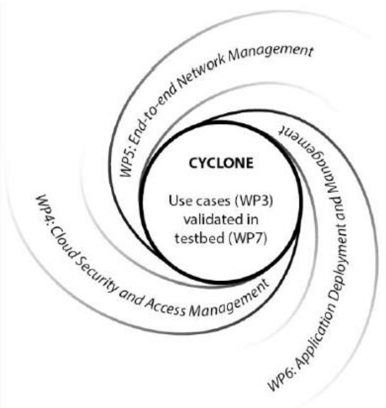

Figure 2: CYCLONE holistic approach and topics overview.

Thus, the holistic cloud management will be loosely coupled and functionally distributed: Loose coupling is achieved through the definition of commonly supported APIs, which can be queried using the holistic cloud model. The functional distribution between all software components makes data exchange and action invocation independent from the communication direction. This distribution makes any combination of integrated software self-contained, i.e., independent from any centralized "cloud manager".

Use Cases

The project has identified two flagship applications: an academic cloud platform and associated services for **bioinformatics research** and a commercial deployment for smart grids in the **energy sector**. These applications will guide the initial development of the tools.

Bioinformatics Use Case

Bioinformatics deals with the collection and efficient analysis of biological data, particularly genomic information from DNA sequences. The capability of modern sequencers to produce terabytes of information coupled with pricing that makes parallel use of many sequencers feasible causes the "data deluge" that is being experienced by researchers in this field.

Bioinformatics software is characterized by a high degree of fragmentation: literally hundreds of different software packages are regularly used for scientific analyses with an incompatible variety of dependencies and a broad range of resource requirements. For this reason, the bioinformatics community has strongly embraced cloud computing with its ability to provide customized execution environments and dynamic resource allocation. The following representative features are being addressed by this CYCLONE use case:

- Securing human biomedical data
- Cloud virtual pipeline for microbial genomes analysis
- Live remote cloud processing of sequencing data

Energy Use Cases

Fulfilling the "202020" climate protection goals of the European Union requires reliance on renewable energy sources that are both decentralized and volatile. The power distribution grid must become smarter to efficiently incorporate such resources efficiently. The Smart Core Interworks platform (SCI-platform) is an advanced data processing platform that integrates modern data processing into the energy management infrastructure. Its diverse features and challenging requirements include:

- Big Data Management
- End-to-End Security
- Autonomy
- Real-Time
- Certification

The cloud-based smart data platform is for complex data processing. The data is collected, stored, analyzed and processed but also made available through different gateways, for different services. The energy use case focuses on Business to Business electrical energy system to improve the energy management infrastructure. The collected and analyzed data does not contain any personal information. The following representative features are being addressed by this CYCLONE use case:

- Deployment of SCI-platform services
- Data retention policies within the storage services
- Anomaly detection within Smart Grids

			Cloud Standards Customer Council					Cloud Computing Use Case Discussion Group			
			CSCC S1	CSCC 2	CSCC 3	CSCC 4	CSCC 5:	CCUC 1	CCUC 2	CCUC 3	CSC-CB
Interoperability	IAAS	Data									
		Service	✓	✓	✓	✓	✓	✓	✓	✓	✓
		Application	✓	✓	✓	✓	✓	✓	✓	✓	✓
		System	✓	✓	✓	✓	✓	✓	✓	✓	✓
	PAAS	Data									
		Service	✓	✓	✓	✓	✓	✓	✓	✓	✓
		Application	✓	✓	✓	✓	✓	✓	✓	✓	✓
		System	✓	✓	✓	✓	✓	✓	✓	✓	✓
	SAAS	Data									
		Service	✓	✓	✓	✓	✓	✓	✓	✓	✓
		Application	✓	✓	✓	✓	✓	✓	✓	✓	✓
		System	✓	✓	✓	✓	✓	✓	✓	✓	✓

CLOUDLIGHTNING project

Overview

CloudLightning proposes to create a new way of provisioning heterogeneous cloud resources to deliver services, specified by the user, using a bespoke service description language. Due to the evolving complexity of modern heterogeneous clouds, a system will be built based on principles of self-management and self-organisation. Service descriptions, provided by prospective cloud consumers, will result in the cloud evolving to deliver the required services. The self-organising behaviour built into, and exhibited by, the cloud infrastructure will result in the formation of a number of potential resource coalitions capable of meeting the service needs. These coalitions will typically be composed of heterogeneous components and thus the quality of service that each could deliver will differ. The user will choose from these offerings and the successful coalition will be commissioned to deliver the service.

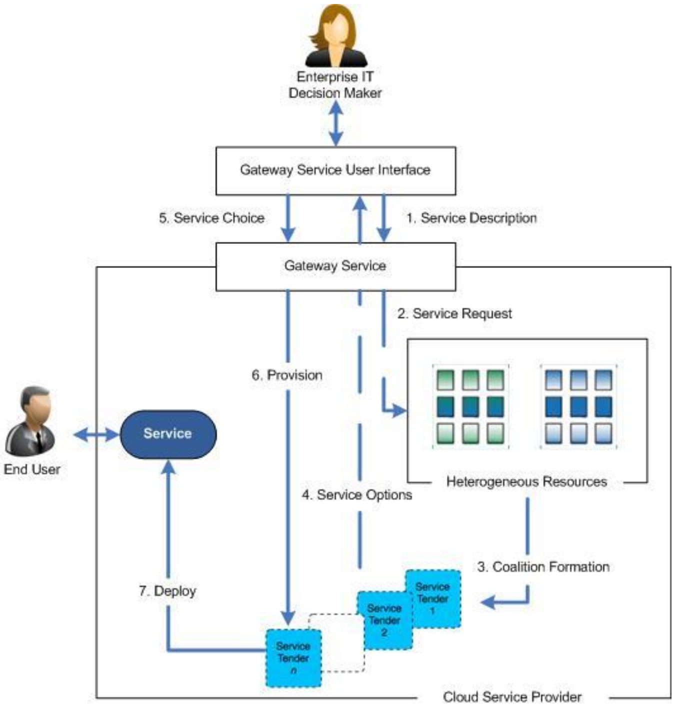

An important objective in creating this system is to remove the burden of low-level service provisioning, optimization and orchestration from the cloud consumer and to vest them in the collective response of the individual resource elements comprising the cloud infrastructure. A related objective is to locate decisions pertaining to resource usage with the individual resource components, where optimal decisions can be made. Currently, successful service delivery relies heavily on the over-provisioning of resources. The CloudLightning goal is to address this inefficient use of resources and consequently to deliver savings to the cloud provider and the cloud consumer in terms of reduced power consumption and improved service delivery, with hyperscale systems particularly in mind.

Approach to Inter-cloud

The CloudLightning approach derives from recognising the importance of restricting the interactions between a Cloud Service Provider and a customer looking from cloud services (labelled, the Enterprise IT Decision Maker in the figure) to a well-defined services interface. Operationally, this customer will declaratively specify a required workflow of services and, in response, the CloudLightning system will deliver access to (possibly a number of alternative heterogeneous) resources capable of delivering these services. A well-defined service interface moves the burden of complex service provision from the customer to the CSP. In this scenario, the customer is responsible only for service description (and SLA negotiations) and the responsibility for service provision, resourcing and delivery moves to the CSP, where the opportunities for providing an optimal solution are greater. Consequently, this model also embodies opportunities for CSPs to collaborate in delivering the individual services specified in a customer's workflow in a manner that hides the complexity of that delivery from the customer, while respecting the customer's SLA.

Use Cases

Use cases from three application domains - Genomics, Oil & Gas Exploration, and Ray Tracing - will be used to validate the CloudLightning management and delivery models. The specific use cases will be augmented with an analysis of dense and sparse matrix analysis techniques that have broad application in a wide variety of fields.

ASCETiC project

Overview

The ASCETiC project approach focuses firstly on the identification of the missing functionalities to support energy efficiency across all cloud layers, and secondly on the definition and integration of explicit measures of energy and ecological requirements into the design and development process for software to be executed on a cloud platform.

ASCETiC's goal is to characterise the factors, which affect energy efficiency in software development, deployment and operations. Its main novel contribution is the incorporation of a novel approach that combines energy-awareness related to cloud environments with the principles of requirements engineering and design modelling for self-adaptive software-intensive systems. This way, the energy efficiency of both cloud infrastructure and software are considered in the cloud service development and operation lifecycle. ASCETiC will demonstrate the outputs of this characterisation through the development of an integrated development environment (IDE) for the deployment of services, the target users being software developers.

ASCETiC project focuses on issues of energy efficient computing, specifically on design, construction, deployment and operation of Cloud services. It proposes novel methods and development tools to support software developers in monitoring and optimizing (minimizing it) the energy consumption resulting from developing and deploying software in Cloud environments.

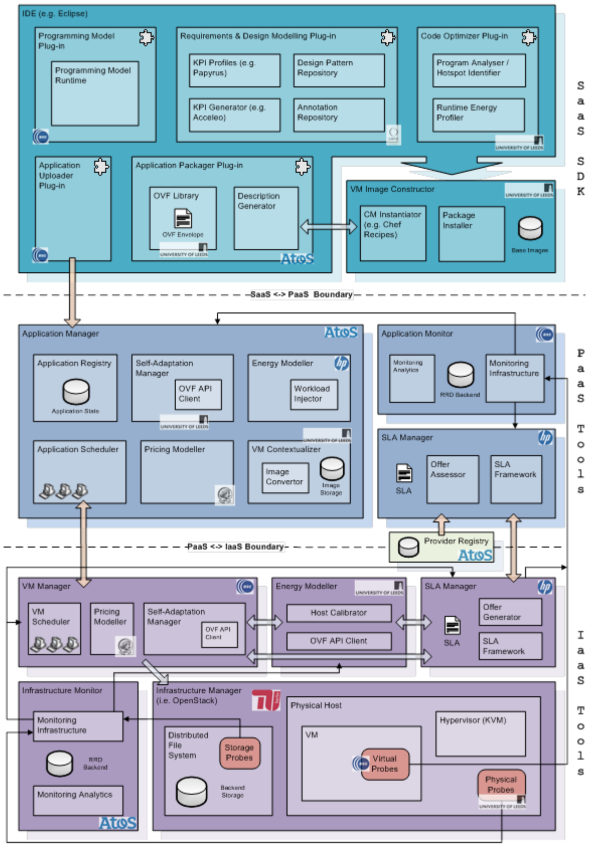

Approach to Inter-cloud

ASCETiC Toolbox functionalities consider three layers of the Cloud stack: SaaS layer facilitates modelling, design and construction of Cloud applications; PaaS layer provides middleware functionality for a Cloud application and facilitates the energy-aware deployment and operation of the application as a whole while IaaS layer considers the admission, allocation and management of virtual resources.

Inter-cloud approach it is implemented at PaaS Layer of the ASCETiC architecture. Platform-as-a-Service layer provides the following features:

- Application deployment management taking into account specific user's deployment requirements.
- Negotiation and management of SLAs for diverse providers making use of energy-aware cost estimation and energy consumption predictions to assess the most suitable provider in which to execute an application.
- Injection of energy measurement probes to the deployment application transparently to the user. Real-time Application Monitoring of resources that are used by a given application.
- Support functions to enable migration of applications among providers, such as VM formats transformation.

In detail, different components of the architecture interact to support these features. Application Manager manages a user's application, which is seen as a set of virtual machine to deploy combined with its deployment requirements. An SLA is negotiated for the deployment with various providers as defined by the Provider Registry. The offers provided by PaaS providers for deployment are assessed by the SLA Manager, which makes use of models generated by the pricing and energy modellers. The applications requirements and the KPIs used to define the conformance are also translated between the IaaS and PaaS layers, in order move energy level metrics up the software stack to the PaaS layer.

Once an agreement has been reached between the PaaS and IaaS layers the application is deployed by the Application Manager to the provider. The VM Contextualizer it is to support portability of image formats and it is used to abstract away issues of cloud deployment environments, which enables focus upon energy efficiency. The Contextualizer also ensures energy probes are setup, which enables the Application Monitor to monitor the application's deployment in multiple providers and confirm that the required energy and performance goals are being achieved by the provider.

Use Cases

			Cloud Standards Customer Council					Cloud Computing Use Case Discussion Group			
			CSCC S1	CSCC 2	CSCC 3	CSCC 4	CSCC 5:	CCUC 1	CCUC 2	CCUC 3	CSC-CB
Interoperability	IAAS	Data									
		Service	✓				✓			✓	(✓)
		Application	✓				✓			✓	(✓)
		System	✓				✓			✓	(✓)
	PAAS	Data									
		Service									
		Application									
		System									
	SAAS	Data									
		Service									
		Application									
		System									
tab	IAA S	Data									
		Service	✓				✓			✓	(✓)

		Application	✓				✓			✓	(✓)
		System	✓				✓			✓	(✓)
	PAAS	Data									
		Service									
		Application									
		System									
	SAAS	Data									
		Service									
		Application									
		System									

ASCETiC addresses CSCC-1 and CSCC-4 by supporting multi provider deployment. This is supported at IaaS level and supporting both inter-operability across providers and VM portability by means of its VM contextualisation tool. Support CUCC-3 in ASCETiC it is done by supporting life VM migration across cloud hosts. CSC-CB, it is at current stage not implemented, although it is in project's plans to support it.

ModaClouds project

Overview

The main goal of MODAClouds project (www.modaclouds.eu) is to provide methods, a deci- sion support system, an open source integrated development environment (IDE) and run-time environment for the high-level design, early prototyping, semi-automatic code generation, and automatic deployment of applications on Multi-Clouds with guaranteed Quality-of-Service (QoS). The concept aligned with this goal was introduced in the early paper [1]. The approach was described in several papers enumerated on the project web site.

The main idea behind MODAClouds is to exploit models and model-driven engineering as the main tools to allow developers and operators of cloud applications to abstract away from the details of the clouds being adopted and to be able to switch from one cloud to the other or to manage applications that are replicated on multiple clouds.

Approach to Inter-cloud

MODAClouds approaches inter-cloud from a cloud user perspective: such users are willing to use different cloud services, even mixing IaaS and PaaS level services, without being burden by the technical differences in the way to exploit such services. For this reason we exploit the term multi-cloud rather than inter-cloud. In the MODAClouds context, multi-cloud is declined in two different ways:

- Applications are built in a cloud agnostic way and can be moved from a cloud to the other without the need for the application owner to know about the details for such clouds.
- If fault tolerance and high availability is very important, MODAClouds support the instantiation of different applications replicas on multiple clouds and manages load balancing and traffic redirection as needed.

As mentioned above, the core of the MODAClouds solution is based on models and model-driven engineering. Such solution is supported by the *MODAClouds Toolbox*. It consists in three main components: 1) Creator4Clouds, an IDE for high-level application design; 2) Venues4Clouds, a decision support system that helps decision makers to identify and select the best execution venue for Cloud applications, by considering technical and business requirements; 3) Energizer4Clouds, a Multi-Cloud run-time environment energized to provide automatic deployment and execution of applications with guaranteed QoS on compatible Clouds.

Use Cases

CSCC S1 and CSCC 2 are the main use cases for our project. We also support CSCC 4 as we offer mechanisms to enable cloud bursting at runtime.

			Cloud Standards Customer Council					**Cloud Computing Use Case Discussion Group**			
			CSCC S1	CSCC 2	CSCC 3	CSCC 4	CSCC 5:	CCUC 1	CCUC 2	CCUC 3	CSC-CB
Interoperability	**IAAS**	Data	✓	✓		✓					
		Service	✓	✓		✓					
		Application	✓	✓		✓					
		System	✓	✓		✓					
	PAAS	Data									
		Service									
		Application									
		System									
	SAAS	Data									
		Service									
		Application									
		System									

SeaClouds project

Overview

SeaClouds project aims to solve the problem caused by the current lack of standardization in cloud services, which pushes cloud customers to end up “locked-in” with the chosen cloud provider(s). In the current situation, it is possible to deploy and monitor a stand-alone application, but not a complex one, and even if frameworks for complex applications on the Cloud can be used, this requires changing the code or using modelling languages.

The project works towards giving organizations the capability of “Agility After Deployment” for cloud-based applications, by supporting developers and application managers through the creation of an open source platform that leverages open standards (such as OASIS CAMP and TOSCA) in order to support the deployment of applications over multiple-clouds, the monitoring of such deployments, and the migration of application modules across different (both public and private) cloud providers if needed.

Approach to Inter-cloud

The SeaClouds Approach is fully compatible with the inter-cloud cluster approach. Actually the SeaClouds approach is based on the concept of multi-cloud service orchestration/management.

The SeaClouds technology is designed to fulfil functional and non-functional properties over the whole application. Applications will be dynamically reconfigured by changing the orchestration of the services when the monitoring detects that such properties are not respected.

So, SeaClouds main objective is the development of a novel platform which performs a seamless adaptive multi-cloud management of service-based applications. More specifically, the objectives of SeaClouds are:

O1) Orchestration and adaptation of services distributed over different cloud providers.

O2) Monitoring and run-time reconfiguration of services distributed over multiple heterogeneous cloud providers.

O3) Providing unified application management of services distributed over different cloud providers.

O4) Compliance with major standards for cloud interoperability.

SeaClouds aims at homogenizing the management of heterogeneous clouds, and at supporting the sound and scalable orchestration of complex software systems across these clouds. Systems developed using SeaClouds will inherently support the evolution of their constituent services, so as to easily cope with needed changes, even during runtime.

To achieve this, SeaClouds employs a user-centric architecture tailored to different aspects of the cloud development life-cycle, providing an open, generic and interoperable foundation from which to orchestrate cloud-based applications.

SeaClouds provides services to monitor and manage cloud providers (both public and private clouds), and leverages service level agreement policies to guarantee the required performance and quality of service across multi-cloud environments.

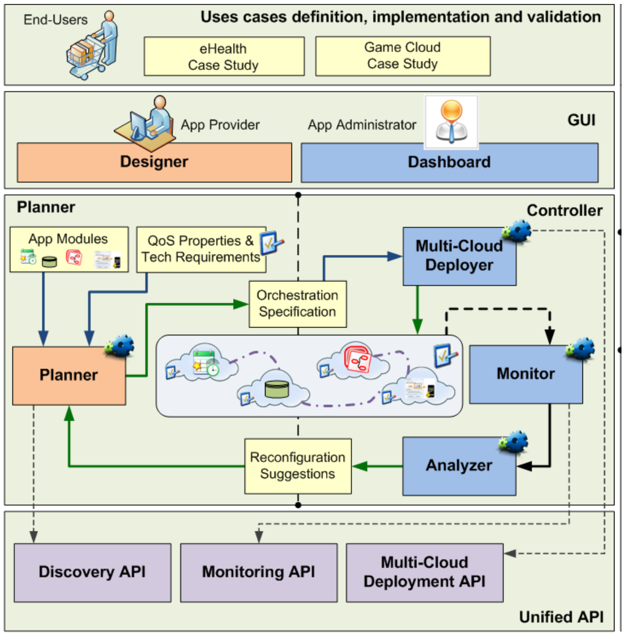

Use Cases

			Cloud Standards Customer Council					Cloud Computing Use Case Discussion Group			
			CSCC S1	CSCC 2	CSCC 3	CSCC 4	CSCC 5:	CCUC 1	CCUC 2	CCUC 3	CSC-CB
Interoperability	IAAS	Data									
		Service	✓	✓		✓					
		Application	✓	✓		✓					
		System									
	PAAS	Data	✓ (partially)	✓ (partially)		✓ (partially)					
		Service	✓	✓		✓					
		Application	✓	✓		✓					
		System									
	SAAS	Data									
		Service									
		Application									

		System									
Portability	IAAS	Data	✓	✓		✓	✓				
		Service	✓	✓		✓	✓				
		Application	✓	✓		✓	✓				
		System									
	PAAS	Data	✓ (partially)	✓ (partially)		✓ (partially)	✓ (partially)				
		Service	✓	✓		✓	✓				
		Application	✓	✓		✓	✓				
		System									
	SAAS	Data									
		Service									
		Application									
		System									

mOSAIC project

Overview

mOSAIC represents a solution for application portability and interoperability of Cloud services over multiple clouds. It aimed to develop an open-source platform that enables application developers to select Cloud services according to their application needs.

Using the Cloud Ontology[21], the Semantic Engine [17] and the Semantic Service Discovery [16], the vendor-agnostic API and various tools, the application developers are able to specify their service requirements and communicate them to the platform.

The selection process is based on the multi-agent brokering performed by the *Cloud Agency*[23] that search for services matching the applications' request. By using mOSAIC approach and software Cloud-application developers and maintainers are able to postpone their decision on the procurement of Cloud services from design time until run-time, while end-user applications are able to find best-fitting Cloud services to their actual needs and efficiently outsource computations and storage.

Approach to Inter-cloud

Using the mOSAIC framework, the decision of which Cloud services from multiple cloud providers to be used is postponed from the design phase to the deployment phase. The key feature is the vendor-agnosticism. The application developer can select at run-time the Cloud services needed from multiple Cloud Providers which can interoperate and possibly be composed in an intercloud scenario.

mOSAIC has developed a new level of abstractions of the Cloud resources that allows a uniform access to multiple Clouds. In particular the Semantic Engine and Dynamic Semantic Discovery Service support the user in discovering the resources and services offered by mOSAIC and multiple Cloud providers, based on Application and Cloud Patterns, and perform their semiautomatic integration in the mOSAIC API. A machine readable (OWL) Cloud ontology has been defined at these purposes, which is being included in the IEEE Intercloud Standard [20]
.

The selection of services from multiple Cloud Providers is semi-automated in mOSAIC by a unique Cloud Agency, a multi-agent systems capable to broker and negotiate the resources and to establish the service-level-agreements with the selected Cloud(s) according to the needs of the applications, and to monitor and possibly dynamically reconfigure the resources provided; six Cloud commercial Cloud providers and six open-source and deployable infrastructure(-as-a-)services are currently semantically represented and connected. The open-source and deployable Platform-as-a-Service is able to manage the selected resources, as well as the application components; particular features are related to the full control of the life-cycle of the application individual components deployed upon infrastructure and software services from multiple Clouds.

Use Cases

			Cloud Standards Customer Council					Cloud Computing Use Case Discussion Group			
			CSCC S1	CSCC 2	CSCC 3	CSCC 4	CSCC 5:	CCUC 1	CCUC 2	CCUC 3	CSC-CB
Interoperability	IAAS	Data									
		Service	✓	✓	✓	✓	✓				✓
		Application	✓	✓	✓	✓	✓				✓
		System	✓	✓	✓	✓	✓				✓
	PAAS	Data									
		Service	✓	✓	✓	✓	✓				
		Application	✓	✓	✓	✓	✓				
		System	✓	✓	✓	✓	✓				
	SAAS	Data									
		Service	✓		✓	✓	✓				
		Application	✓		✓	✓	✓				
		System	✓		✓	✓	✓				
Portability	IAAS	Data									
		Service	✓	✓			✓			✓	✓
		Application	✓	✓			✓			✓	✓
		System	✓	✓			✓			✓	✓
	PAAS	Data									
		Service	✓	✓			✓		✓		
		Application	✓	✓			✓		✓		
		System	✓	✓			✓		✓		
	SAAS	Data									
		Service	✓				✓	✓			
		Application	✓				✓	✓			
		System	✓				✓	✓			

CSCC 1 (Customer Switches Providers for a Cloud Service) and CSCC 5 (Migration of Customer Capabilities into Cloud Services)

The Semantic Engine main aims is to support the user in selecting Cloud APIs components and functionalities needed for building the application on the cloud, and the list of needed resources to be acquired from the Cloud providers. Thus the approach of the Semantic Engine can be used in the context of the use case scenario CSCC 5. In fact through the use of the Semantic Engine it's possible to describe semantically the application by using an application pattern. This pattern can be mapped on proper cloud patterns that fulfil application requirements [19]. Using Cloud Patterns can ease the migration of an in-house complex application to the Cloud and enhance interoperability among different platforms because they make easier to understand the exact functionalities and responsibilities of a specific Cloud application component, which can be at a later time be substituted with a compliant one having the same or similar characteristics. In a scenario in which the application is built in an in-house system the Semantic Engine suggests the components to use to deploy the application on the Cloud, while if the application is already on the cloud the engine will suggest, by using inference rules and the implemented discovery mechanisms, the components from other providers useful to replace the components currently used.

This aspect is also related to the use case scenario CSCC 1 because using semantically annotated Cloud Patterns, users can exploit the available knowledge base, containing meaningful relationships between Patterns' participants and possible implementing Cloud Services, to automatically retrieve the Cloud components needed to migrate their applications across Cloud platforms or from on-premises infrastructure to new Cloud environments [18].

Also, the use of services offered by multiple providers would be automatically suggested and the information necessary for their interoperation (data and API formats, security protocols and so on) would be provided accordingly. This covers the interoperability aspects of CSCC 1 e CSCC 5 at service, application and system levels.

In conjunction with the semantic engine, the mOSAIC API and the Cloud Agency make possible the actual migration by offering agnostic API for uniform access to services from different Cloud Provides.

CSCC 2 (Customer Uses Cloud Services from Multiple Providers)

The mosaic project, and in particular the mOSAIC platform and the Cloud Agency, fully meets the needs of the use case scenario CSCC 2. In particular, the mOSAIC platform-agnostic application programming interface enable to use multi-Cloud resources at PAAS level in a transparent way through the use of agnostic connectors, while the Cloud Agency, by using different vendor agents is able to acquire IAAS resources from different cloud providers and manages them with a unique interface.

CSCC 3 (Customer Links One Cloud Service to Another Cloud Service) and CSCC 4 (Customer Links In-house Capabilities with Cloud Services)

The problem of make interoperable different cloud services (or in-house capability with cloud services) to implement new functionalities addressed by the use case scenarios CSCC 3 and CSCC 4 can be solved by using the semantic support of the mOSAIC project. In particular the Semantic Engine is able to discover not only services but also a composition of Cloud Services that could be used to implement a more complex functionality. Moreover the use of the Dynamic Discovery Service is useful to discover similarities and differences among cloud APIs and thus to point out possible lack of partial functionalities that can be solved through the combination of different cloud services.

CCUC 1 (Changing SaaS Vendors)

In the context of the use case scenario CCUC 1 the mOSAIC Semantic Engine is once again capable of recognizing the equivalent services able to replace the service in use by supporting common formats.

CCUC 2 (Changing middleware vendors)

The mOSAIC platform offers an API, implemented in Python, Java, and Erlang, to develop components which run on top of its platform. mOSAIC's API is designed to be event-driven, and the communication among mOSAIC components takes place through message queues. The mOSAIC's basic component is the Cloudlet, an event-driven and stateless component whose functionalities do not depend on the number of its instances at run-time (has a degree of autonomy). The Cloudlets are able to access cloud services through Connectors. The concept of Connector is introduced to ensure the independence from the cloud service interfaces. A Connector is a concrete class that abstracts the access to cloud resources and defines the set of events to which the Cloudlet should react. The Connectors access cloud services using Drivers, which actually implement the cloud services interfaces. They can be interpreted as wrappers of native resource APIs or uniform APIs. Until now mOSAIC' software repository includes modules for more than ten Public Clouds. Among these, mOSAIC supports well known providers like Amazon, Rackspace, and GoGrid, as well as European Cloud providers including Flexiant (UK), CloudSigma (Switzerland), NIIFI (Hungary), Arctur and Hostko (Slovenia), latest two using VMwares vCloud. Moreover, Private Clouds are built by using open-source technologies like Eucalyptus, OpenNebula, CloudStack, and the already cited OpenStack. This covers the CCUC 2 use case scenario, because by using the mOSAIC API it is possible to change cloud middleware vendors in a transparent way.

CCUC 3 (Changing VM hosts) and CSC-CB (Cloud Bursting)

Multi Agent Systems represents a meaningful support in multi cloud environment due to their ability to automatize operation such as brokering, negotiation, monitoring and reconfiguration.

Using the Cloud Agency it's possible to monitor the resources on which the application is deployed through mobile software agents which take measures inside the Cloud resources, take decision based on the monitoring values or using automatic settings and eventually decide to migrate the application on other IaaS provider. During the migration phase it's possible to take benefit of the multi agent system to manage with an agnostic interface resources of different providers (through vendor Agents which implement wrapper for specific Clouds), negotiate with other providers and reconfigure the application on the new chosen cloud provider [22]. This meets the needs generated by CCUC 3 and CSC-CB by providing an agnostic interface able to manage IAAS resources of multiple cloud providers.

SWITCH project

Overview

The SWITCH project (Software Workbench for Interactive, Time Critical and Highly self-adaptive cloud applications) addresses the urgent industrial need to be able to develop and execute time critical applications in cloud environments.

Many time critical applications (persistent applications that must adapt rapidly to time-sensitive events) often have very high business value (e.g. on-demand business collaboration platforms) or social impact (e.g. disaster early warning systems). These applications demand a high standard of Quality of Service (e.g. minimising tsunami emergency response time) or quality of experience (e.g. ensuring smooth delivery of ultra-high definition audio and video for live events), but are very difficult to develop and operate because of their distributed nature and the high requirements they impose on the runtime environment—in particular because of the sophisticated optimization mechanisms needed to develop and integrate system components that must interact seamlessly with one another.

Cloud environments are capable of providing, on demand, the virtualized, elastic, controllable and high quality services needed to support these kinds of complex, distributed applications. Indeed many cloud providers already provide most of the technologies needed to develop and deploy these applications. However, what time critical applications still need from the cloud is the ability to control the selection and configuration of infrastructural components in response to changing requirements and environmental pressures. Unfortunately current Cloud environments lack the tools and application programming interfaces that would allow the developers to exert such control on the underlying infrastructure in an intelligent, semi-autonomous manner.

SWITCH provides an interactive and flexible software workbench that can provide the necessary tools to control the lifecycle for rapid development, deployment, management and dynamic reconfiguration of complex distributed time-critical Cloud applications.

Approach to Inter-cloud

The core idea of the SWITCH environment, an application-infrastructure co-programming and control model, will be developed for time-critical cloud applications, and is envisaged to support inter-cloud as well as single cloud scenarios. The new model brings together application composition, execution environment customisation, and runtime control, which are normally treated as separate processes, into one optimisation loop based on adherence to time critical constraints. In this model: 1) the application logic will be specified along with any QoS/QoE requirements, together with the required programmability and controllability of the underlying cloud environment. Existing formats such as TOSCA could be of help in this context, perhaps the TOSCA specification which could be extended with the SDN - related part in the course of the SWITCH project. 2) A virtual runtime environment will be customised to meet the critical application requirements, and be provisioned in the cloud with Service Level Agreements (SLAs) suited to the application. The negotiations for resources with various supported Clouds (including SLA) must include the networking part. And 3) The application will autonomously adapt its own behaviour and that of the virtual environment when performance drops at runtime. To realise this model, SWITCH requires an abstraction layer for both applications and infrastructure, with abstract semantic descriptions for both executable components (which can be made portable via containerisation) and programmable infrastructure (exploiting technologies such as SDN for cloud interoperation). The SWITCH environment employs formal performance reasoning mechanisms to guide each step in the development, and tools are delivered to the users via three subsystems:

- The *SWITCH Interactive Development Environment* (SIDE) subsystem provides interfaces for all user- and programmer-facing tools, by exposing a collection of graphical interfaces and APIs that tie SWITCH's services to a Web-based environment.
- The *Dynamic Real-time Infrastructure Planner* (DRIP) subsystem prepares the execution of the applications developed in the SIDE subsystem by: 1) modelling semantically the different QoS/QoE attributes; 2) defining an optimal virtual runtime environment (possibly across multiple clouds); 3) creating (and negotiating) a Service Level Agreement with the resource providers; and 4) deploying the platform required by the application.
- The *Autonomous System Adaptation Platform* (ASAP): 1) monitors the status of the application and the runtime environment; 2) checks adherence to the required quality attributes; 3) autonomously manipulates the application and runtime environment to maintain optimal system level performance against time critical constraints; and 4) learns from its own decision history to improve future decisions.

Use Cases

			Cloud Standards Customer Council					**Cloud Computing Use Case Discussion Group**			
			CSCC S1	CSCC 2	CSCC 3	CSCC 4	CSCC 5:	CCUC 1	CCUC 2	CCUC 3	CSC-CB
Interoperability	**IAAS**	Data	✓	✓	✓	✓	✓			✓	✓
		Service	✓	✓	✓	✓	✓			✓	✓
		Application	✓	✓	✓	✓	✓			✓	✓
		System	✓	✓	✓	✓	✓			✓	✓
	PAAS	Data	✓	✓	✓	✓	✓			✓	✓
		Service	✓	✓	✓	✓	✓			✓	✓
		Application	✓	✓	✓	✓	✓			✓	✓
		System	✓	✓	✓	✓	✓			✓	✓
	SAAS	Data									
		Service									
		Application									
		System									
Portability	**IAAS**	Data	✓	✓	✓	✓	✓			✓	✓
		Service	✓	✓	✓	✓	✓			✓	✓
		Application	✓	✓	✓	✓	✓			✓	✓
		System	✓	✓	✓	✓	✓			✓	✓
	PAAS	Data	✓	✓	✓	✓	✓			✓	✓
		Service	✓	✓	✓	✓	✓			✓	✓
		Application	✓	✓	✓	✓	✓			✓	✓
		System	✓	✓	✓	✓	✓			✓	✓
	SAAS	Data	✓	✓	✓	✓	✓			✓	✓
		Service	✓	✓	✓	✓	✓			✓	✓
		Application	✓	✓	✓	✓	✓			✓	✓
		System	✓	✓	✓	✓	✓			✓	✓

AppHub project

The aim of the AppHub project is to support the market outreach strategies of EU-supported open source by launching AppHub, the European open source market place. AppHub is a service platform that will help the market to seamlessly identify, position and implement the software outcomes of these projects. The partners that will develop, run and promote AppHub over this two-year project and beyond combine unparalleled expertise in open source community management, EU research projects and a breakthrough technology in software asset management.

AppHub will be based on three interrelated services:

- The AppHub Directory allows placing software assets as part of a reference architecture and thus identifying rapidly ways to compose various open source assets into a service architecture.
- The AppHub Factory lets users build and maintain full software stacks as templates using a visual "point and click" interface or APIs.
- The AppHub Marketplace provides users with self-service access to pre-packaged business and IT applications via a customizable, white-labelled app store, and to deploy them in various cloud infrastructures.

AppHub supports Inter-cloud Challenges, Expectations and Issues Cluster project´s to:

- to improve the quality of the software they produce.
- offering a viable market place that allows those projects to showcase their software and to disseminate their results in a “hands-on” approach.
- Provide a technical infrastructure for the rapid deployment of open source software into arbitrary cloud infrastructures.

4. Conclusions

This document has presented the initial work of the Inter-cloud Challenges, Expectations and Issues Cluster with regards to identification of Research Areas and Challenges and Projects classification.

Presented Work in Research Areas and Challenges aims to be useful inputs to consultation process for future work programmes for the areas of Cloud Computing and Software, reflecting vision brought from existing projects and collaborating institutions in Inter-Cloud related projects. Next steps in this cluster will consider updating this initial work so to refine and further develop this analysis in a more detailed roadmap for future research.

In addition, the work on analysis of EU funded research projects in which we have classified existing project's work in relation to Inter-operability approach, Topics and Scenarios addressed, Standards used and enabled Use Cases allows the Inter-Cloud cluster to have a baseline to be used to achieve its ambitions for collaboration among participating projects. To this end, the cluster has stated to identify a structure of Technical Working Groups, in order to facilitate specific technical collaborations and to deliver Technical collaborations across projects. So far the following identified Technical working groups have been identified: Semantics, Model-Driven, Portability and Network. Additional work will follow to identify more working groups, as well as, across cluster collaboration with standards, specifically for this cluster in the scope of TOSCA.

5. References

[1] Petcu, D. (2013). Multi-Cloud: Expectations and Current Approaches. In Proceedings of the 2013 international workshop on Multi-cloud applications and federated clouds - MultiCloud '13 (p. 1). New York, New York, USA: ACM Press. doi:10.1145/2462326.2462328

[2] Nikolay Grozev and Rajkumar Buyya. 2014. Inter-Cloud architectures and application brokering: taxonomy and survey. Softw. Pract. Exp. 44, 3 (March 2014), 369–390. DOI:http://dx.doi.org/10.1002/spe.2168

[3] Gartner Survey Analysis: Cloud Adoption Across Vertical Industries Exhibits More Similarities Than Differences, https://www.gartner.com/doc/2987617/survey-analysis-cloud-adoption-vertical

[4] Ferrer, A. J., Hernández, F., Tordsson, J., Elmroth, E., Ali-Eldin, A., Zsigri, C., … Sheridan, C. (2012). OPTIMIS: A Holistic Approach to Cloud Service Provisioning. Future Gener. Comput. Syst., 28(1), 66–77. doi:10.1016/j.future.2011.05.022

[5] Toosi, A. N., Calheiros, R. N., & Buyya, R. (2014). Interconnected Cloud Computing Environments. ACM Computing Surveys, 47(1), 1–47. doi:10.1145/2593512

[6] Cloud Standards Coordination - Final Report, http://ec.europa.eu/digital-agenda/en/news/cloud-standards-coordination-final-report

[7] Topology and Orchestration Specification for Cloud Applications Version 1.0, http://docs.oasis-open.org/tosca/TOSCA/v1.0/cs01/TOSCA-v1.0-cs01.pdf

[8] Topology and Orchestration Specification for Cloud Applications (TOSCA) Primer Version 1.0, http://docs.oasis-open.org/tosca/tosca-primer/v1.0/cnd01/tosca-primer-v1.0-cnd01.pdf

[9] Cloud Application Management for Platforms Version 1.1, http://docs.oasis-open.org/camp/camp-spec/v1.1/camp-spec-v1.1.html

[10] Beniamino Di Martino, Giuseppina Cretella, and Antonio Esposito. 2015. Advances in Applications Portability and Services Interoperability among Multiple Clouds. IEEE Cloud Comput. 2, 2 (March 2015), 22–28. DOI:http://dx.doi.org/10.1109/MCC.2015.38

[11] Jclouds, https://jclouds.apache.org/

[12] Deltacloud, https://deltacloud.apache.org/

[13] Libcloud, https://libcloud.apache.org/

[14] Baudoin, C., Dekel, E., Edwards, M. et al.: Cloud Standards Customer Council. Interoperability and Portability for Cloud Computing: A Guide (2014). http://www.cloudstandardscustomercouncil.org/CSCC-Cloud-Interoperability-and-Portability.pdf

[15] Ahronovitz, M., Amrhein, D., Anderson, P. et al.: Cloud Computing Use Cases-white paper. http://www.cloud-council.org/Cloud_Computing_Use_Cases_Whitepaper-4_0.pdf

[16] Cretella, G., & Di Martino, B. (2013, December). Semantic and matchmaking technologies for discovering, mapping and aligning cloud providers's services. In Proceedings of International Conference on Information Integration and Web-based Applications & Services p. 380-384. ACM, 2013

[17] Cretella, Giuseppina, and Beniamino Di Martino. "A semantic engine for porting applications to the cloud and among clouds." Software: Practice and Experience (2014).

[18] Di Martino, B.; Esposito, A.; Cretella, G., "Semantic Representation of Cloud Patterns and Services with Automated Reasoning to support Cloud Application Portability," in Cloud Computing, IEEE Transactions on , doi: 10.1109/TCC.2015.2433259

[19] DI MARTINO, Beniamino; CRETELLA, Giuseppina; ESPOSITO, Antonio. Mapping design patterns to cloud patterns to support application portability: a preliminary study. In: Proceedings of the 12th ACM International Conference on Computing Frontiers p. 50-58. ACM, 2015.

[20] Di Martino, B.; Cretella, G.; Esposito, A.; Willner, A.; Alloush, A.; Bernstein, D.; Vij, D.; Weinman, J., "Towards an Ontology-Based Intercloud Resource Catalogue -- The IEEE P2302 Intercloud Approach for a Semantic Resource Exchange," in Cloud Engineering (IC2E), 2015 IEEE International Conference on , vol., no., pp.458-464, 9-13 March 2015

[21] MOSCATO, Francesco, et al. An analysis of mosaic ontology for cloud resources annotation. In: Computer Science and Information Systems (FedCSIS), 2011 Federated Conference on. IEEE, p. 973-980, 2011.

[22] VENTICINQUE, Salvatore; TASQUIER, Luca; MARTINO, Beniamino Di. A restfull interface for scalable agents based cloud services. International Journal of Ad Hoc and Ubiquitous Computing, p. 219-231, 2014.

[23] VENTICINQUE, Salvatore; TASQUIER, Luca; MARTINO, Beniamino Di. Agents based cloud computing interface for resource provisioning and management. In: Complex, Intelligent and Software Intensive Systems (CISIS), 2012 Sixth International Conference on. IEEE, p. 249-256, 2012.

[24] Rick McGeer, Distributed Computing, http://ec.europa.eu/information_society/newsroom/cf/dae/document.cfm?action=display&doc_id=6768

[25] Massimo Re Ferre, vCloud Architect, “Common Challenges When Moving to the Cloud and How to Avoid Them”, VMware vCloud Blog, April 2013, https://blogs.vmware.com/vcloud/tag/reliability

[26] Cloud Standards Coordination - Final Report, http://ec.europa.eu/digital-agenda/en/news/cloud-standards-coordination-final-report