Future Cloud

Federated Cloud Reference Architecture Position Paper

February 2021

Participant projects and representatives

Name	Affiliation	Project
Beniamino Di Martino	University of Campania	mOSAIC
George Kousiouris	Harokopio University of Athens	PHYSICS
Pedro Garcia Lopez	Universitat Rovira i Virgili	CloudButton
Philippe Massonet	CETIC	BEACON, SPARTA
Sébastien Dupont	CETIC	BEACON
Patrizio Dazzi	CNR	ACCORDION
Daniel Vladušič	XLAB	SODALITE
Ana Juan Ferrer (Editor)	Atos	PLEDGER, DECENTER
Leire Orue-Echevarria	TECNALIA	DECIDE, PIACERE, MEDINA
Demetris Trihinas	University of Nicosia/University of Cyprus	RAINBOW
Zhiming Zhao	University of Amsterdam	SWITCH, ARTICONF
Zoltan Mann	University of Duisburg-Essen	FogProtect
Sebastian Scholze	ATB	SmartCLIDE
Enric Pages	Atos	Elastest

Contents

1. Introduction

The aim of Future Cloud Cluster is to provide a forum for discussion and collaboration for research and innovation initiatives that address next generation Cloud Computing challenges and issues, including diverse forms of distributed computing (Cloud, Multi-Cloud, Edge, Fog, Ad-hoc and Mobile computing).

These comprises research on diverse areas encountered in Future Cloud Infrastructure and software management such as: the integration of heterogeneous devices ranging from high performance computing to Edge devices in computing infrastructures; the development and design of decentralised service-oriented systems; improvements to virtualization, data and workload abstraction technologies; service management and QoS; or the automation the organisation and management of the back-end resources.

In addition, the cluster studies research in the interoperability and portability of across these diverse and heterogeneous computing environments.

Moreover, we address research topics and trends associated with the hyper distribution of computing; These include the consideration of interoperability, portability, elasticity, self-organisation across many and heterogeneous resources in edge clouds, private enterprise clouds, aggregated cloud models and large Cloud set-ups; the orchestration and placement problems that address heterogeneity and trade-offs among consistency, reliability, availability, cost and performance.; and management models that cope with decentralisation and distribution of computing considering enhanced workload portability abilities, richer fault-tolerance verification with complexities at level of extreme scalability and parallelisation.

Today a number of research projects are already analysing these research areas from diverse perspectives and approaches. The main aim for this cluster is to foster collaboration among these existing research initiatives. Table 1 Cluster Participant Projects provides the list of existing EU funded projects that participate in the Future Cloud Cluster initiative.

This document elaborates on the Future Cloud Cluster inputs to the European Commission with our vision on Cloud Federation Reference Architecture, its rationale and comparison with existing Federation models. With this main purpose, we update the Future Cloud Cluster Vision for Future Cloud.

Relying on it, the proposed Reference architecture in Section 4, Reference Architecture aims at bringing a concrete materialisation of this vision. The Future Cloud reference architecture develops at the Cloud Federation, Cloud and Edge computing layers. For each of these layers, it elaborates on the necessary features and building blocks. In addition, we identify H2020 Future Cloud Cluster research projects that have provided or will provide research findings in that can contribute to the realisation of these building blocks. To finalise, we associate Future Cloud Research Areas to our reference architecture building blocks.

Table 1 Cluster Participant Projects

Project	Name	URL
ACCORDION	Adaptative Edge/Cloud compute and network to support nextgen applications	https://www.accordion-project.eu/
AppHub	The European Open Source Market Place	http://www.apphub.eu.com
ASCETiC	Adapting Service lifecycle towards Efficient Clouds	www.ascetic-project.eu
BASMATI	Cloud Brokerage Across Borders For Mobile Users And Applications	http://www.basmati.cloud
BEACON	Enabling Federated Cloud Networking	www.beacon-project.eu
Cloudbutton	Serverless Data Analytics Platform	https://cloudbutton.eu/
CLOUDLIGHTNING	Self-organising, Self-managing heterogeneous cloud	www.cloudlightning.eu

Project	Name	URL
CloudPerfect	Find YOUR perfect Cloud	cloudperfect.eu
CloudSocket	Business and IT-Cloud Alignment using a Smart Socket	http://www.cloudsocket.eu
CYCLONE	Complete Dynamic Multi-cloud Application Management	www.cyclone-project.eu
DECIDE	DevOps for trusted, portable and interoperable multi-Cloud applications towards the Digital single market	www.decide-h2020.eu
DITAS	Data-intensive applications Improvement by moving daTA and computation in mixed cloud/fog environmentS	www.ditas-project.eu
Elastest	An elastic platform to ease end to end testing	elastest.eu
ENTICE	Decentralised repositories for transparent and efficient virtual machine operations	www.entice-project.eu
EUBrasilCloudForum	Fostering and International Dialogue between Europe and Brazil	eubrasilcloudforum.eu
FED4IOT	Federating IoT and cloud infrastructures to provide scalable and interoperable Smart Cities applications, by introducing novel IoT virtualization technologies	https://fed4iot.org/
FogProtect	Protecting Sensitive Data in the Computing Continuum	https://fogprotect.eu/
H-CLOUD	The Forum for Strategy Focused Cloud Stakeholders	http://www.h-cloud.eu/
INPUT	INPUT Project	input-project.eu
IOSTACK	Software-defined Storage toolkit for Big Data	iostack.eu
mF2C	Towards an Open, Secure, Decentralised and Coordinated Fog-to-Cloud Management Ecosystem	www.mf2c-project.eu
MICHELANGELO	MICHELANGELO Project	www.mikelangelo-project.eu
ModaClouds	Model-Driven Approach for design and execution of applications on multiple Clouds	www.modaclouds.eu
mOSAIC	Open-Source API and Platform for Multiple Clouds	www.mosaic-cloud.eu
M-SEC	Multi-layered Security technologies to ensure hyper connected smart cities with Blockchain, BigData, Cloud and IoT	http://www.msecproject.eu
PaaSage	Model-based Cloud Platform Upperware	www.paasage.eu
PLEDGER	Paving the way for Next-generation Edge computing	http://www.pledger-project.eu
PHYSICS	oPtimized HYbrid space tIme Service Continuum in faaS	https://cordis.europa.eu/project/id/101017047
RAPID	Heterogeneous Secure Multi-Level Remote Acceleration Servies for Low-Power Integrated System and Devices	rapid-project.eu
RAINBOW	An open, trusted fog computing platform facilitating the deployment, orchestration and management of scalable, heterogeneous and secure IoT services and cross-cloud apps	https://rainbow-h2020.eu/
RECAP	Reliable Capacity Provisioning and Enhanced Remediation for Distributed Cloud Applications	recap-project.eu
SeaClouds	Agility after the deployment	www.seaclouds-project.eu
SmartCLIDE	The Stairway to Cloud	https://smartclide.eu/
SODALITE	SOftware Defined AppLication Infrastructures managemenT and Engineering	https://www.sodalite.eu/
SPECS	Secure Provisioning of Cloud Services based on SLA Management	http://www.specs-project.eu/
SSICLOPS	Scalable and Secure Infrastructures for Cloud Operations	www.ssiclops.eu
SWITCH	Software Workbench for Interactive, Time Critical and Highly self-adaptive Cloud applications	www.switchproject.eu

2. Future Cloud Cluster Vision for Future Cloud

It is widely recognised that Cloud computing has allowed the democratisation of computing. The development of IaaS, PaaS and SaaS models[1] witnessed over the last decade, has permitted to transition from "an era in which underlying computing resources were both scarce and expensive to an era in which the same resources started to be cheap and abundant"[2]. Public Cloud offerings from hyperscale providers today deliver the illusion of infinite compute resources and make reality the radical commoditization of computing envisioned by utility computing models. While offerings of Cloud Hyperscalers concentrate at present times the majority of the user demand for public Cloud computing, new alternative models, such as Cloud Federation[3], start to emerge with the purpose of facilitating interoperability and portability across providers.

The concept of Federated Cloud has been an area extensively studied in the Cloud research community. Federated Cloud models that encompass Hybrid Cloud and Multi-clouds are today gaining momentum both in commercial and community set-ups. Examples of these are Hybrid Cloud Solutions in the market, from vendors such as vmware[4] and more recently AWS Outposts[5] , IBM/RedHat OpenShift or Google Anthos among others.

Hybrid Cloud is the simplest form of a Federated Cloud scenario. In this scenario, an established Cloud installation (classically a private one) is able to shift workloads to another existing Cloud provider (often a public one). This model is the enabler for Cloud bursting scenarios, in which theoretically providers are able to move without service interruptions applications between two cloud installations or providers when demand for computing exceeds the capacities of the original provider.

In more advanced federated cloud scenarios, a cloud provider is capable of acquiring additional capacity from a set of other providers. In addition, it can offer available capacity to an established cloud federation. From a user perspective, parts of a service are placed in a combination of resources which span the original provider and the selected set of remote providers. This enables enhanced elasticity, resilience and fault tolerance. Typically, the original cloud provider remains as the contact point for the user to observe the agreed QoS. The federated cloud scenario is often related to community cloud set-ups. The notable examples in the eScience community in Europe is represented by the European Open Science Cloud (EOSC)[6].

From a commercial perspective, Federated Cloud can act as a model for cloud providers that own multiple cloud islands (for instance, in diverse geographical regions) with the purpose of balancing workload among them. But more importantly, it can enable groups of providers to offer communally services to final users, including or not a mediation broker. This is the model envisioned by GAIA-X[7].

In order to realize a fully Federated Cloud vision, additional aspects need to be developed. These encompass areas such as: collective provisioning, metering and billing; across cloud privacy, security and identity management; fine grained QoS and composable Service Level agreements; secure mechanisms for data sharing; consideration of diversity and heterogeneity of resources at all levels (compute, data and network); as well as, adoption of existing Cloud standards. It is solely with the complete development of these novel capabilities when interoperability among cloud providers will become a reality, raising new business opportunities both for existing and new cloud stakeholders.

An additional element of the development of the fully Federated Cloud Vision are standards. The existing standards, such as the standards coming from DMTF, ETSI, ITU, OASIS, etc. must be considered, used or even contributed to, to create a common understanding between different components of the Federated Cloud. Of course, de-facto standards, used by the global market leaders, must be incorporated, to make the Federated Cloud indeed useful. To make this a reality, an effort in planning and management of the overall Federated Cloud idea is required, with the significant expected benefits in simplification of connectivity between different components and providers themselves.

Federated Cloud is expected to allow the exploitation of hybrid cloud models to their highest potential in a rich Cloud ecosystem. A sustainable Federated Cloud market will allow enterprises and public sector to wholly materialise the Cloud promise of a complete hybrid provisioning which achieves the right balance among services provisioned in house and those consumed from external Cloud providers. From providers' perspective, Federated Cloud models seek to overcome the traditional disjunctive between public and private cloud, hence, to include the necessary services across the full cloud continuum, embracing set-ups for large data storage, HPC intelligent analytics and Edge for IoT.

In addition, instantiations of cloud marketplaces which take into consideration operational concerns and technological challenges applicable to a specific sector will enable vertical markets customizations. In these, federated cloud vertical marketplaces will need to combine general purpose services, secure mechanisms for data sharing and monetisation, as well as, specific services tailored for the needs of vertical sectors.

Whilst the major investments into the already developed technologies should not be abandoned, new technical developments are required. Here, we should rely on the pre-existing well-adopted open-source software. The core premise of the approach is to not reinvent the wheel, but rather use the foundations and build on them, with the EU requirements in mind and the using the EU approach towards security and privacy. We stress that the way to widespread adoption of the Federated Cloud approach is dependent on the familiarity with the tools and approaches in SMEs, which does not require major retraining of the workforce.

While these Federated Cloud markets are today being materialised, the emergence of Edge computing even broadens the concept. Edge computing is today rapidly growing and represents the major trend in the distributed computing area. Edge computing, also named Fog computing by some authors, is capitalising the attention of both research and commercial communities. Edge computing represents the first step towards the decentralisation of Cloud computing, bringing the concept of Federated Cloud to its next evolutionary stage.

Edge computing develops the idea of providing novel forms of computing embracing computing power and data resources increasingly obtainable at the Edge of the network. Edge computing forces existing Cloud computing environments, which emerged as part of a centralisation paradigm, to evolve to decentralised environments avoiding drawbacks of large data movements and latency, specifically found in IoT scenarios.

The call towards the decentralisation of Cloud computing is present in a wide variety of works and under diverse terms. Satyanarayanan[8] and Lopez[9] contextualise the current trend towards Cloud computing decentralisation in the context of alternating waves of centralisation and decentralisation, that have affected computing since the '60s. In these, centralisation of computing has been prevalent in '60s and '70s through batch processing and time sharing, and in the 2000s employing traditional centralised Cloud computing models; whereas alternating with decentralisation in '80s and '90s via the emergence of personal computing and in which Edge and its extended Cloud Federation models presents the last episode of this on-going trend towards hyper-distribution of computing.

By means of the development of Edge computing, the concept of Federated Cloud is transformed in order to encompass an even more distributed approach which includes a diversity of cloud and edge offerings that lead to better performance which enable a wider diversity of application and services, complementarity to traditional cloud X-as-a-service models.

At the same time Edge computing is itself evolving in many diverse areas. Many of today's conceptions of Edge computing take the view of a single Edge device located in proximity to an IoT installation. Even in this simple scenario, many challenges still remain in relation to the optimal workload encapsulation, service placement, networking, security and privacy. When this scenario is complemented with the necessary orchestration of Edge resources with the rest of the surrounding compute continuum, additional challenges emerge.

Edge computing developments regard the Edge as stationary single device environments which provide computing and storage services to a set of nearby IoT devices. In these, commonly, IoT devices are only viewed as sources of data, and their increasing sophistication in terms of computing and storage capacity is ignored. However, the materialization of IoT is instrumental for the incorporation of compute resources not only on dedicated non-mobile Edge computing devices, but more and more also on a wide diversity of non-traditional heterogeneous devices distributed at the Edge of the network. The exploitation of the capacity in these innovative devices which are often mobile, bring the need to develop novel methods for managing resource dynamicity, churn and scale.

At the same time, the rising data availability resulting from IoT deployments together with the recent advances and popularity of AI and Deep learning at the Edge is expected to increase computing demand at the Edge by several orders of magnitude. This calls for the development of novel data management mechanisms and software tools for the compute continuum which permit intelligent value extraction from the huge volumes of generated data at the Edge of the network.

Besides, novel compute architectures, such as neuromorphic computing (brain inspired computing), are currently aiming at surpassing von Neumann's architectures by designing compute units which target energy optimal AI workloads execution at the Edge. This complements an existing source evolution that Cloud providers are already taking up, by means of exploitation of hardware heterogeneity. Developing computing continuum technology compatible with the management of hardware heterogeneity requires to find new ways to optimally endeavour heterogeneous special purpose processing units without losing the advantages of abstraction of the utility-based models. These could be materialised by the development of: novel virtualisation tools, heterogeneity-aware scheduling at resource and platform levels, in addition to programming models that permit to handle heterogeneity seamlessly. Moreover, taking advantage of hardware heterogeneity (together with specific developments for energy aware metering and scheduling) are key aspects to enable the advance towards energy aware, optimised and sustainable cloud continuum compute services.

Furthermore, the movement towards hyper distribution of the computing continuum will have to allow to develop federated cloud and edge computing scenarios that inherently contemplate intelligent collaboration, self-organisation, self-management and self-healing across many and heterogeneous resources present in all kinds of IoT Edge devices, micro edge data centres, private enterprise clouds, federated cloud models and large Cloud set-ups. Orchestration and placement problems in this context will necessitate of novel management tools, programming models and approaches potentially p2p and bio-inspired that cope with heterogeneity and hyper distribution of computing while handling security and performance together with extreme scale, parallelisation and fault-tolerance.

3. EU Cloud Federation rationale

This section presents the Future Cloud Cluster assessment on benefits and barriers for adoption of European wide federations. In addition, it presents Use cases from the Research projects which participate in the cluster activities stressing the identified benefits which can be extracted from Edge to Cloud federation approaches. To finalise, it studies existing Cloud federation initiatives and exhibits a comparison among their different approaches.

3.1 Benefits

For the providers

- Benefit from the long tail business model: while each CSP may have several big clients for whom the infrastructure is set up, smaller companies may benefit from the unused, but available, infrastructural resources, providing the CSP with smaller, but continuous, income.
- Continuous innovation to remain competitive.
- Implementation of compute continuum approaches.

For the consumers:

- More offerings to select based on their own requirements, for instance, cost, availability, a certain 'legal' level
- If the federation follows a rigorous approach to endorse services in the catalogue of federated services, the consumer can feel that their data and applications are more secure.
- Optimized management of cloud services
- Operational efficiency is improved
- Enablement of "best venue" selection depending on application requirements.

3.2 Barriers

Some of the barriers for federation include:

- Federation is based on voluntary basis
- Different levels of coupling in federations
- Lack of trust because the on-boarding process is not open and it is not sure that the services on-board follow existing accreditations and have a compliance check.
- Existing investments
- Interoperability and portability aspects among different services are not ensured.
- Perceived security and privacy risks.

3.3 Use cases

This section presents Use cases that exemplify benefits to gain from Edge to Cloud continuum federation in diverse sectors and end user applications from existing H2020 Cluster research projects.

Use case	***Collaborative VR***
H2020 Research project	ACCORDION
Use case description	
OVR, a startup focused on simulation for medical training, provides a gamified multi-user VR platform built on Unity engine that exploits the MAGES SDK on top of a networking layer by Unity. The MAGES SDK handles and synchronizes in-game interactions, deformable object transformations and physics	

simulation, broadcasts transformation values over the network while the custom Geometric Algebra interpolation engine supports interpolation of in-between positions/rotations locally at each end-device. Currently, for the case of untethered Head-Mounted Displays (HMDs) the storage, rendering, compression are all local processes as part of a single application component that is installed and run on the untethered HMD. The transitioning to the ACCORDION platform will support the OVR business case to develop and promote collaborative cloud VR training applications specially formulated for untethered HMDs and the adaptation of OVR's networking layer to edge computing will optimize the current status of the cooperative mode, ensuring lower latency and higher performance on average network conditions and ultimately a higher number of CCUs.

In order to enable this transitioning and leverage the edge-cloud architecture for the given application, computation offloading has to be exploited to support migrating part of its computing in edge-cloud resources, requiring interactions and data exchanges between the different modules placed on device, edge and cloud, considering the benefit of remote execution, the cost of data transmission and the complexity of application partitioning. As part of the computation offloading paradigm, and important for collaborative multi-player environments, is a networking solution that effectively manages communication among clients and handles the game state with less dependency on the master server. Offloading this functionality from the client-host and providing it as a data service from ACCORDION to the application can be realized via a relay server that is able to handle beyond the broadcasting of messages, also host migration and solve for game-state continuity.

BENEFITS GAINED FROM EDGE / CLOUD FEDERATION

- Computation offloading to automatically migrate part of its computing in federated edge-cloud resources located near-by
- Execution of part of the application on clouds and part at the edge by taking into account the benefit of remote execution, the cost of data transmission and the complexity of application partitioning

USE CASE	***MULTI PLAYER MOBILE GAMING***
H2020 RESEARCH PROJECT	ACCORDION

USE CASE DESCRIPTION

ORBK is a company that produces and sells multiplayer mobile games. Game servers will be deployed on top of the ACCORDION system to meet the requirements of NextGen mobile gaming
, which aims to lower latency between servers and clients and highly improve user experience.

BENEFITS GAINED FROM EDGE / CLOUD FEDERATION

- To take advantage of AI-based network orchestration to dynamically and automatically deploy new servers based on performance metrics and player's geographical localization.

USE CASE	***QOE OPTIMIZATION IN CONTENT DELIVERY***
H2020 RESEARCH PROJECT	ACCORDION

USE CASE DESCRIPTION

Plexus is a technology company that produces Traqus. It is a geo-localization platform which processes and analyzes the information obtained from mobile devices detected by access points within a WiFi network. Anblick provides digital signage and works as a content manager where content can be programmed to reproduce on screens, for example in a hospital, train station or mall. Currently these cloud solutions are independent, and information is not exchanged between them.

The transitioning to the ACCORDION platform will support the Plexus business case to provide personalized content by integrating both products incorporating edge computing to ensure lower latency and higher performance on average network conditions.

The approach towards ACCORDION raises the development of a brand-new complete platform, based on a PaaS model integrating edge computing and derived services. A multi-level, microservices-based scenario platform that calculates varied information from different locations and creates real-time analytics. Analytics and derived knowledge elements that provide advantages to all parts of a new value chain based on information processing.

BENEFITS GAINED FROM EDGE / CLOUD FEDERATION

− ACCORDION platform allows processing of device information on the edge, reducing latency and allowing for advanced calculations to be performed in the cloud. Based on the results of these calculations adapted content will be streamed to the different screens within the network. − The continuous processing of information will be managed from different leveled nodes. The integration of all the information will lead to an expert system whose parts are mutually feeding each other. As a result, content delivery to different devices and the measurement of the QoE of responses will be done in a consistent way across peripheral devices and connections.

USE CASE	***METABOLOMICS SERVERLESS DATA ANALYTICS***
H2020 RESEARCH PROJECT	CloudButton
USE CASE DESCRIPTION	
EMBL (European Molecular Biology Laboratory) has developed three metabolomic data analytics pipelines using Serverless technologies (Lithops) from H2020 CloudButton project. In particular, Lithops (lithops.cloud) is a Python framework enabling simple deployment of python code to many Cloud and Edge backends. The technology is used by the Metaspace.eu community. One pipeline is explicitly benefitting from Multi-cloud/Hybrid Cloud technologies that enables to combine EMBL in-house resources, Mestaspace participant´s resources, and public Clouds.	
BENEFITS GAINED FROM EDGE / CLOUD FEDERATION	
− Cost-aware optimizations of leveraging private computing resources (EMBL) and combine them in a transparent way with federated multi-cloud resources − Simplicity/Productivity for data scientists which only learn local Python libraries, and their code is transparently optimized to Edge/Cloud Federated resources.	

USE CASE	***CLOUD AND EDGE NETWORK FEDERATION FOR AUTOMOTIVE PLATOONING***
H2020 RESEARCH PROJECT	BEACON and SPARTA
USE CASE DESCRIPTION	
Platooning involves grouping cars on the highway under a platoon leader to optimise traffic and energy consumption. The platoon leader is under the supervision of a traffic management platform that communicates zone policies to the platoon leader. The traffic management platform is deployed in a cloud and communicates with the platoon leaders via edge platforms deployed along the road system. The cloud platform and the edge platforms are federated into a cloud network federation that can be managed with a coherent global network policy including a security policy. The network federation could be extended to the platoon network itself, i.e. the network that is setup between the platoon leader and the individual vehicles of the platoon (V2V communications).	
BENEFITS GAINED FROM EDGE / CLOUD FEDERATION	
● Cloud/edge network QoS can be managed globally for the Cloud/edge federation, and adapted to changing communication needs. ● A global security policy can be defined and enforced including monitoring functionalities to take into account the safety critical nature of a platoon. This includes security monitoring and response automation functionalities.	

USE CASE	***MIXED REALITY APPLICATIONS***
H2020 RESEARCH PROJECT	PLEDGER
USE CASE DESCRIPTION	
Enhance the capabilities of the AR/VR/MR solutions by coupling them with edge computing technologies and the corresponding load allocation and optimisation tools, in order to provide high level services in industrial environment. Head-mounted displays for industrial use have the disadvantage of limited memory and processing power. At the same time, computation capacity for achieving high resolution graphics at the edge is imposing new requirements for performance, scalability and quality of service: features that bring in a necessary trade-off capability between what computation has to be implemented on the edge (even through GPU processing) and what needs to be transmitted for further processing to the cloud infrastructure.	

Benefits gained FROM Edge / Cloud Federation
− Faster and more powerful visualisation of 3D CAD models, display 3D models of even higher performance in real world environments enabled y compute continuum exploitation. − Secure handling of confidential data at Edge.

Use case	***Edge infrastructure for enhancing safety vulnerable road users***
H2020 Research project	PLEDGER
Use case description	
The goal of this use case is to assist drivers with enhanced perception capabilities, in order to reduce the number of accidents, especially with Vulnerable Road Users . In particular, it will be focusing on intersections where the drivers' visibility may be obstructed, thus missing the presence of a pedestrian or a bicycle approaching. Through smart Edge Cloud infrastructure, the pedestrian or the bicycle will be detected and the surrounding drivers will be notified about their presence.	
Benefits gained FROM Edge / Cloud Federation	
− Exploitation of diverse edge computing configurations (device and infrastructure) connect over the last mile network and work in harmony to reduce collisions, improve traffic, and increase life safety. − Enhanced QoS by reducing low latency and jitter, high availability and security to prevent tracking and disclose of privacy sensitive data.	

Use case	***Human-Robot Collaboration in Industrial Ecosystems***
H2020 Research project	RAINBOW
Use case description	
In today's large industrial ecosystems, the production process demands the involvement of humans and robots to assembly heavy and complex entities (e.g., cars, engines), with robots carrying the heavy objects. In such settings, the prominent safety factor for consideration is the prevention of human-robot collisions. Thus, real-time indoor localization services are deployed to monitor the flow of objects and detect human worker positioning collaborating with machinery. However, indoor positioning for safety-critical industrial IoT requires the propagation of telemetry and positioning data at millisecond range from thousands of objects, human workers and robotics via (mobile) sensors. At the same time, it requires the execution of complex AI algorithmic models on 3D-spacing topologies to output coordination plans and prevent collisions among humans and robotic machinery.	
Benefits gained FROM Edge / Cloud Federation	
− Because of the delay-sensitive nature of collision detection services, edge infrastructure is employed to process -in place- positioning data across the computing continuum. − Provisioning of backup fog nodes to ensure high-availability of collision detection services. − High-scalability by ensuring that no centralized cloud offering is overwhelmed with data and queuing compute tasks.	

Use case	***Manufacturing the data mining on edge***
H2020 Research project	PLEDGER
Use case description	
Industrial Internet of Things platform to connect to all machines in a factory and extends them with intelligent algorithms that increase both profitability and production quality. Data recording in combination with self-learning algorithms allow the IIoT platform to precisely determine the condition of a machine, predict any pending maintenance, and ensure complete component traceability (partly hosted on the edge and in the cloud). These results are processed in individually tailored dashboards for the operator and enable both an overview and detailed analyses of the production and machines.	
Benefits gained FROM Edge / Cloud Federation	

Distributing data and computations across the edge and cloud might have the following benefits: Rich data sets for better machine learning: For machine learning applications, a large dataset with high quality is essential. High performance computing: Edge devices have limited computation power. Running software in the cloud allows to perform computationally intense analysis or machine learning. Security: The edge solution of cybernetics is isolated from the internet and runs locally close industrial to machines

USE CASE	***DIGITAL TRANSFORMATION OF URBAN MOBILITY***
H2020 RESEARCH PROJECT	RAINBOW
USE CASE DESCRIPTION	
Establishing real-time geo-referenced notification systems for vehicles traveling in urban areas about critical situations for the city mobility network, due to any possible cause (e.g., severe weather, failure of road infrastructure, huge congestion) is a daunting challenge. A geo-referenced vehicle notification system requires enabling existing sources of information on city traffic into sharing modality (e.g., traffic lights, vehicles, sensors), identifying new sources of information (e.g., crowdsourced citizens' reports) and securely processing and distributing -in time- alerts to all the relevant users, indicating the level of criticality via mobile apps and vehicle navigation software.	
BENEFITS GAINED FROM EDGE / CLOUD FEDERATION	
− The optimal “splitting” of computations between the on-board application, the edge and the cloud backend. − The identification of the geographical location of heterogeneous MEC servers to support different user populations and densities, tailoring the service to the peculiarities of the covered area. − Dynamic infrastructure adaptation (edge, MEC, cloud) to highly mobile and dynamic loads in geo-distributed environments.	

USE CASE	***POWER LINE SURVEILLANCE VIA SWARM OF DRONES***
H2020 RESEARCH PROJECT	RAINBOW
USE CASE DESCRIPTION	
Power line inspection using helicopters and ground patrols is expensive, time consuming and dangerous, especially when dealing with difficult terrains, large geographic areas, and unpredictable weather. The deployment of drone swarms presents a golden opportunity for monitoring critical infrastructures. Drones can follow precise flights along power lines and maintain high overlapping of imaging strips. To achieve this, flight planning is integrated with localization of power pylons so that a routing plan with mission waypoints is derived prior departure. However, drone technology is not without challenges. The foremost challenge is drone autonomicity. Performing high quality imaging is energy consuming, resulting in the frequent return of drones for recharging. In turn, image analysis is performed offline after drones return to base without any indication if the images are sufficient. If not, the task must be repeated again.	
BENEFITS GAINED FROM EDGE / CLOUD FEDERATION	
– Drone energy consumption optimization via low-cost adaptive monitoring by dynamically adjusting the drone data collection intensity and communication rate based on data positioning and current images. – Secure and real-time flight routing optimization upon applying advanced data management on top of a trusted overlay mesh network connecting the drone swarm. There is no need to move data from the ‘trusted’ network for processing.	

USE CASE	***SMART PRECISION AGRICULTURE***
H2020 RESEARCH PROJECT	PHYSICS
USE CASE DESCRIPTION	
Greenhouses are nowadays the most sophisticated way to control plant environment to increase their production, reduce impact of climate uncertainty, provide physical barriers to diseases, enabling strong reduction of chemical pesticides. However, they require more and more parameters to be set by the grower (e.g., 200 in a standard soil-less glasshouse used for tomatoes). As a consequence, parameters	

are mostly set to default values, without adaptation to the location of the farm, the needs of the species and of the cultivar, their potential in yield and quality (dry matter and sugar content). Thus a more dynamic, online process should be pursued in order to gather collected data, model and constantly optimize parameter setting in the greenhouse.
Benefits gained FROM Edge / Cloud Federation
– To ensure robustness to computational and data breakout, a system is needed with multiple computation locations. A complete version of the model resolving the problem in external locations (Cloud or more powerful edge nodes) for the entire greenhouse (and for multiple greenhouses) and a lighter version in the greenhouse for results application – Enable the easier adaptation of the traditional application graph in the FaaS model as well as the process from an offline sporadic one to a continuous monitoring and operation function – Optimize running costs, automate distribution and management of components and application performance.

<table>
<tr><td><b>Use case</b></td><td><b><i>eHealth Personalized Monitoring and Collective Analysis</i></b></td></tr>
<tr><td><b>H2020 Research project</b></td><td>PHYSICS</td></tr>
<tr><td colspan="2"><b>Use case description</b></td></tr>
<tr><td colspan="2">People with mild conditions, more serious chronic ones, or in rehabilitation, are being remotely monitored and advised in their common, every-day settings by medical professionals. Professionals of para-medical domains (wellness, fitness, nutrition) monitor and coach their clientele. People themselves monitor their quality of life, keeping track of adverse events like migraines or athletic injuries and the effect of those on their lifestyle. Such systems are most important in the post-emergency COVID-19 stage the world is now entering, where large portions of the population need to be monitored for symptoms, incurring the smallest possible strain to the health systems. In such cases the authorities and healthcare institutions need to monitor and advise the entire population while not in the national health system, but at home.All the above situations have in common the (self) monitoring of people for long periods of time in their everyday life, for collecting and analysing real-world data (RWD) in order to coach them. The process of acquiring data from a multitude of diverse devices is challenging both in terms of i) interfaces, protocols etc , ii) regarding data harmonization, transformation and storage, iii) privacy/security preservation. Furthermore, the transformation to a FaaS oriented version of the analytics and machine learning algorithms used in the analysis is in many cases complicated and cannot reach the scale needed for a more general use of the platform e.g. in cases of pandemics or large scale clinical trials.</td></tr>
<tr><td colspan="2"><b>Benefits gained FROM Edge / Cloud Federation</b></td></tr>
<tr><td colspan="2">– Enable dynamic design and adaptation of the monitoring logic and functional incorporation , enhancing the ability to define adaptation logic for each type of device, while the achieved semantic elevation will enable the incorporation of relevant adapters, privacy preserving and security enhancing function nodes for data transmission in the application workflow
– Significantly improve running costs of the platform, especially in cases when further analysis of the data needs to be performed in an event driven manner and based on the monitoring data (e.g. intensification of monitoring). This is highly coupled with the event driven nature of FaaS as well as with the ability to investigate provider performance and associated costs. New algorithms based on the cloud design patterns can more efficiently and rapidly be generated as well as dynamically deployed across the continuum, including location constraints due to legal or IPR related issues</td></tr>
</table>

<table>
<tr><td><b>Use case</b></td><td><b><i>Smart Manufacturing for increased resilience and interplay</i></b></td></tr>
<tr><td><b>H2020 Research project</b></td><td>PHYSICS</td></tr>
<tr><td colspan="2"><b>Use case description</b></td></tr>
<tr><td colspan="2">Typically production lines are comprised of a variety of different h/w devices, each with completely different characteristics and capabilities.The challenge of utilizing all these diverse setups in a seamless and integrated manner while optimizing their assignment and ensuring their functional incorporation is enormous. Aspects that need to be addressed include adding real time deployment and redundancy features, while making use of 3rd party knowledge and services for a variety of functional aspects, as well as divide logical and physical deployment.The use case will demonstrate how to transform</td></tr>
</table>

manufacturing architectures to serverless architecture, it will make use of the FaaS principle and applied semantics to enhance the flexibility and adaptability of modular production lines.

BENEFITS GAINED FROM EDGE / CLOUD FEDERATION

- Semantic description of h/w capabilities of diverse devices and hardware types in the industrial environment, as well as their comparison and matchmaking to application components needs and requirements
- Increased resilience in production lines during component deployment as well as optimized usage of the available resources, while expressing requirements such as locality of computation, latency etc, enhancing the real-time reaction of the system to unexpected changes/failures.
- Enable easier incorporation and interplay between cloud, edge and other available resources for meeting operational real-time targets, deciding on the distribution of necessary functions to available or external resources

3.4 Cloud Federation related initiatives

In this section we explore existing initiatives such as Gaia-X, NIST Federation Reference Architecture and International Data Spaces which highlight diverse existing approaches for Cloud Federation.

GAIA-X

GAIA-X aims at defining and implementing an “Infrastructure and Data Ecosystem according to European values and standards” and its goal is to ensure the digital and data sovereignty of Europe [10].

The high level and conceptual architecture of GAIA-X, represented as an X, shows in the upper layer all aspects related to data governance (i.e. policies and usage control, interoperability),and trust (e.g. federated identity and access management), while in the bottom layer the infrastructural aspects are shown. This includes the federated catalogue of services that to be on-boarded in GAIA-X need to be described in a certain way, following the ontology defined for that purpose and the requirement to comply with existing certification schemes, such as BSI C5 or the upcoming European cloud security certification scheme. This layer also includes the continuous metering of the service, monitoring of the compliance of the service with respect to SLA and the continuous auditing or continuous compliance with the accreditations mentioned when being on-boarded.

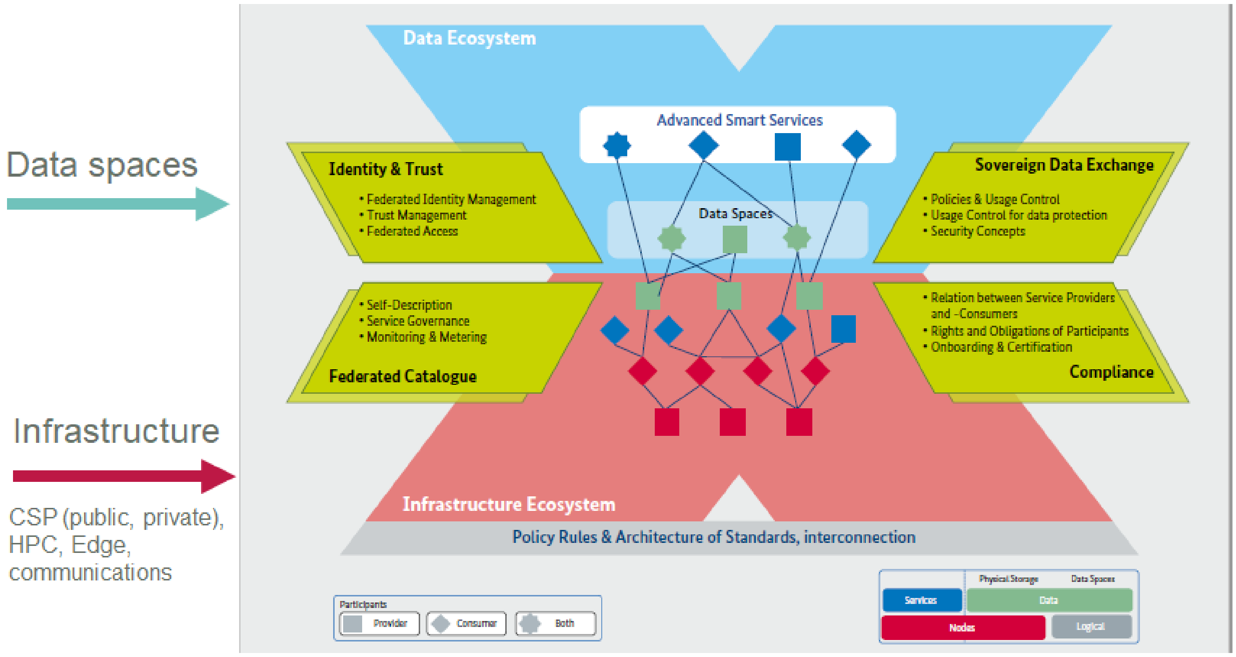

The core architecture elements of GAIA-X include the concepts of Node, Service, Service instance or data asset, described as follows:

- Node: a computational resource. Hierarchies are supported to accommodate those service providers that have multiple locations for the offering of their services.
- Service: a service is a cloud offering
- Service Instance is the realization of a Service on Nodes
- Data Asset: a data set that is made available to Consumers via a Service that exposes the Data Asset.

Figure 1 GAIA-X conceptual architecture (source GAIA-X)

GAIA-X foresees different integration means of the data, services and nodes. This integration can be at a horizontal level, where various services coming from different providers can work together, and at a vertical level, covering the stack of node - service and alternatively, data.

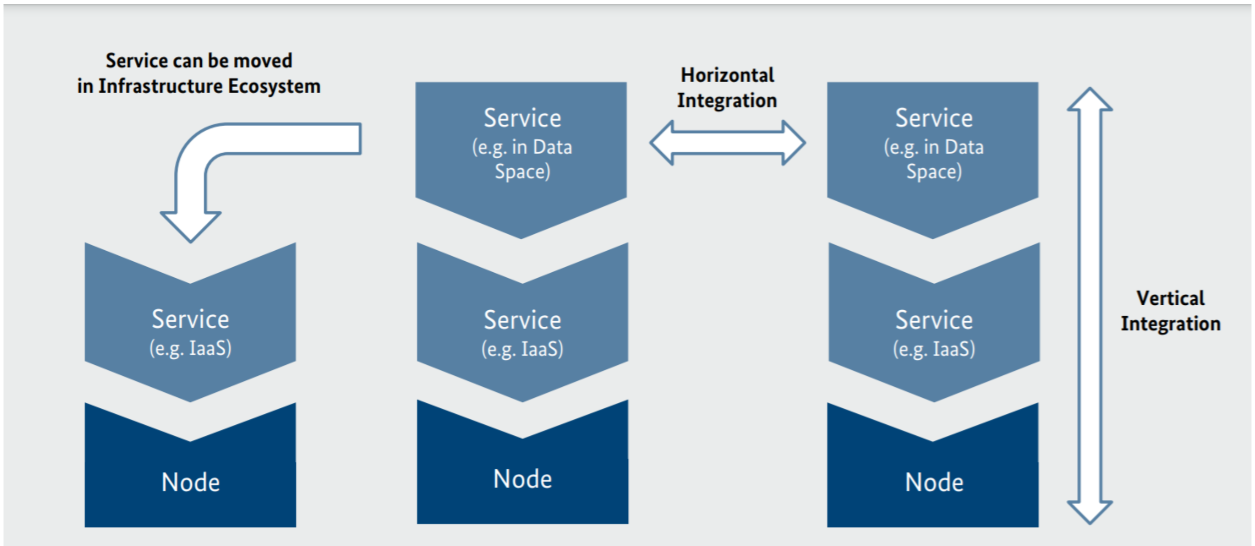

Figure 2 Possible horizontal and vertical service integration in GAIA-X (source: GAIA-X)

In GAIA-X the concept of federation addresses the following challenges (source: Gaia-x):

- Decentralized processing locations
- Multiple technology stacks
- Regulated market and policy requirements
- Multiple stakeholders

Several aspects that are very important in GAIA-X are as follows:

- Identity and trust, offered at both levels, which include Federated identity, federated access and trust management among the different GAIA-X interactions possible
- Interoperability, both at data an service level, to ensure a seamless data exchange and favour also portability among the different service providers
- Control and monitoring: in the case of data exchange, the policies and the use of data that ensure an effective data protection are key. In the case of the infrastructural services, the monitoring comes

at different stages: monitoring of the services offered, and continuous monitoring of the certifications accredited by the provider

- Self-description of a node, and service, following a defined ontology.
- Security

NIST Federation Reference Architecture

US National Institute of Standards and Technology (NIST) presented in February 2020 the NIST Cloud Federated Reference Architecture (CFRA)[11]. This architecture model focuses on the service deployment and service orchestration aspects in the context of Community Clouds.

CFRA Model defines the essential characteristics of a Cloud Federation relying on the concepts of Virtual Organisations in Grid computing. These characteristics comprise four main elements:

- Virtual Administrative Domains, "*A federation is a virtual security and collaboration context that is not necessarily "owned" by any one user or organization".*
- Membership and Identity Credentials: "Since only specific users, sites, and organizations collaborate for common goals, these participating entities have membership in the federation and identity credentials that are linked to each member. "
- Shared Resource Metadata and Discovery: *"Users, sites, and organizations can participate in a federation by choosing to share some of their resources and metadata and making them discoverable and accessible to other federation members."*
- Governance: *"Participating members agree upon the common goals and governance of their federation, based on well-known roles, attributes and policies."*

Figure 3 presents CFRA Actors. These represent major stakeholders in Cloud Federation together with main functions identified for each actor.

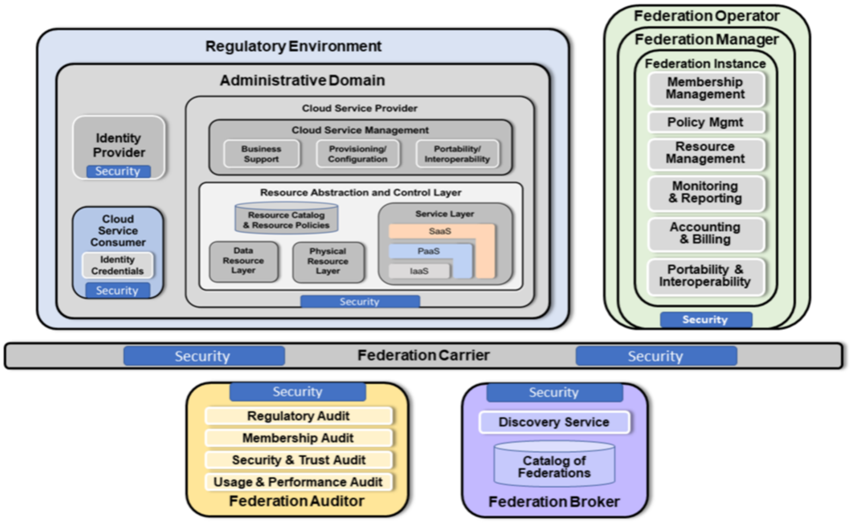

Figure 3 The NIST Cloud Federation Reference Architecture Actors (Source CFRA)[12]

Industrial Data Spaces

The International Data Spaces (IDS)[13] initiative objective is to standardize data exchange and sharing among participant entities while ensuring the sovereign usage of data. Its mission revolves around four main elements:

- Secure data exchange in which the data owner always keeps control over its data use.

– Developing an international standard for data governance, architecture and interfaces
– Based on well-defined use cases
– Acting as the basis in which to provide a variety of software solutions, smart services and business models.

Recently in “GAIA-X and IDS Position Paper”[14] , a combination of Gaia-X architectures and IDS has been explored. This exploration concludes on the complementarity of both initiatives for the realisation of end-to-end data value chains. This is facilitated by one side, by means of the IDS enablement of data spaces for smart services in industry verticals achieving data sovereignty: and Gaia-X focus on cloud and infrastructure sovereignty, on the other side.

3.5 Comparison of Federation Models

Up to the present, the concept of Cloud Federation has addressed, as its primary concern, how to enable the orchestration of workloads among diverse public or private cloud installations. In this context multiple terms have been employed to name cloud federation models.

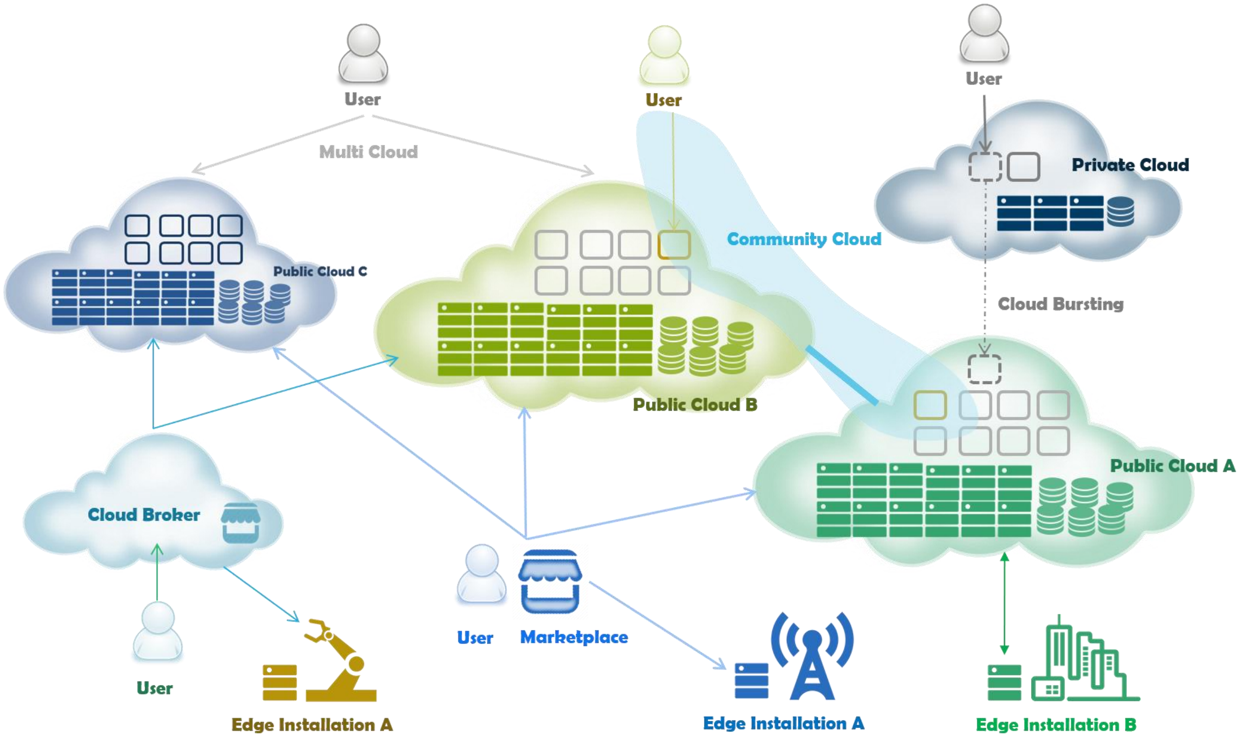

Figure 4 Representation of Cloud Federation models

Terms such as Multi-cloud, Cloud-brokerage and Cloud Bursting directly relate to workload management across cloud installations, in which Federation decision is taken by a centralised actor. These can be classified according to two main criteria: decision-actor and decision-time.

❑ Decision-actor determines whether the decision to use diverse clouds is taken in a centralised manner by a cloud provider, by a cloud broker or event directly by the cloud user.
❑ The decision-time classifies scenarios considering that the cloud federation decision occurs at the time of deploying an application, or it, on the contrary, happens once the application is in execution in a provider selected provider.

Fully Decentralised Cloud Federation decision models, such as those envisaged in the Future Cloud Vision in Section 2 Future Cloud Cluster Vision for Future Cloud, build their routes in research works, such as Swarm Computing[15], and are not yet present in the previously analysed Federation initiatives. Figure 4 and the following table represents the identified different interactions in Cloud Federation models.

Table 2 Comparison Comparison of Cloud Federation Models per Decision Actor and Time

		Decision Actor		
		Centralised Cloud User	**Centralised Cloud Broker**	**Centralised Cloud Provider Private/ Public / Community**
Decision time	**Deployment**	Multi-Cloud	Multi-Cloud Cloud Brokerage Cloud Marketplace	Community Cloud
	Execution	Hybrid Cloud		Cloud Bursting Hybrid Cloud
Examples		*VMware HCX*[16] *Google Anthos*[17]	*GaiaX*	*NIST Federated Cloud* *EGI*[18]

These cloud federation models typically address cloud IaaS or PaaS stack levels. New initiatives, such as the previously introduced IDS and Gaia-X develop the new concept of Data federations and marketplaces. These Data federations concepts have the purpose of enabling trustworthy and sustainable data sharing schemes among diverse federation participants. These new types of federations, therefore, develop new criteria for classifying federations:

- ❑ Cloud Federation at Workload level are those Cloud federations that have the goal of sharing at level of services and application workloads.
- ❑ Cloud Federations at Data level, build on top of Cloud Federations at workload level, introducing the ability to share data on top of services and applications.

Federations considering the Data level are related to the new ability to bring data up to the cloud enabled by the addition into the Cloud Federation models of IoT and Edge computing installations. As represented in Figure 4, in current commercial deployments, Edge software stacks are specifically bounded to a specific proprietary set of Cloud services addressing both data and storage management. Extended Cloud Federation concepts enlarging Cloud Federation concept to Edge, as such present in Gaia-X and IDS, facilitate the data access, exchange, and monetisation.

Another important factor to consider within Cloud Federations comparisons is the level of resource visibility across federation members. Commercial Cloud Federation models, both a public, private Cloud and Edge levels, always consider API based access. In these models, public APIs are available to the federation to handle workloads in the providers, however the exact configuration of available compute and data resources made available to the federation, as well as, their internal toolset utilised for its management, is unknown to the rest of the federation. These approaches hence develop loosely coupled Federations, in which each of the participants do not need to give visibility to the rest of the federation of the exact configurations and toolsets, and remain in full control on how they handle the capacity in their infrastructure.

Differently, the Cloud Federation model proposed by NIST as well as EGI Cloud Federation model[19] REF[d] present tightly coupled federations that rely on a Community Cloud approach developed as analogy to Grid computing Virtual Organisation. These models also rely on APIs to operate within the federation and have their routes in eScience approaches for collaboration among research infrastructures. This model assumes a set of resources being made available to the federation in its associated virtual organisation. Consequently, the set of resources being made available to the infrastructure are handled separately from the rest of the providers' infrastructure. At the level of the toolset for management, for instance, EGI determines a certain toolset to be employed for infrastructure management, offering options based on OpenStack[20] and Open Nebula[21].

Table 3 Comparison of Cloud Federation Models per Federation level and approach

	Federation Level			**Federation Approach**	
	Edge+Cloud	**Cloud(IaaS/PaaS)**	**Data**	**Loosely coupled**	**Tightly coupled**
Examples	*GaiaX*	*NIST Federated Cloud* *VMware HCX*[22] *Google Anthos*[23]	*Gaia-X* *IDS*	*Gaia-X* *Cloudera*	*NIST Federated Cloud* *EGI*[24]

4. Reference Architecture

In this section we present a proposal from the Future Cloud Cluster for a Reference Architecture for Cloud federation. The architecture is aligned to the Future Cloud Vision characteristics and features detailed in Section 2. For this architecture, we start by presenting its main building blocks elaborated in three different levels: Federation management, Cloud and Edge computing. In section 4.2we present briefly Future Cloud Cluster research projects that are elaborating / have elaborated components and research findings of interest for building the features intended per each level and building block. To finalise this section, we associate research areas developed in our H2020 Future Cloud Research Roadmap to each one of the building blocks defined in this Reference architecture.

4.1 FutureCloud Reference Architecture

The Future Cloud reference architecture is developed in three different levels: Federation management, Cloud and Edge computing. These are presented in detail in the upcoming subsections.

Federation Management

As presented in the introduction the concept of Federated Cloud is conceived as the natural evolution of Cloud and IoT, while at the same time, being a combination of hybrid cloud architectures with Edge computing.

It is important to note that federation set-ups will be managed using the principle of subsidiarity of intelligence, meaning that decision making will happen at the lowest appropriate level, in which the different cooperating environments are autonomous. Intelligence and knowledge about the overall status of the Federation is decentralised and spread across diverse participating instances.

At the same time the economics of participation will be managed to capitalise individuals' and communities' willingness to engage, identify and deliver assets. Time-constrained reservation, adaptive selection, conflict resolution and techniques to consider the volatility and uncertainty (introduced by real-world dynamics) will be developed to enable efficient and reliable service provision.

A key objective, but also a challenge for such federated systems, is breaking the interoperability barriers and fences which characterize the technology vendor ecosystems fostering a heterogeneous integration with standardized technologies and pluggable systems. It is important to note that security is a crucial element for the implementation of Federated Cloud Environments.

In the federation concept, different infrastructures, services and systems need to be made interoperable and seamlessly exploitable. Such heterogeneous environment is provided by an ecosystem of cooperating stakeholders that aim to share their own resources for sustainability or business purposes. From this point of view, it is important that the federation promotes a collaborative and not competitive coexistence of stakeholders, mapping agreements in the federation with actual resource and service exploitation and guaranteeing fair opportunities. This implies investigating solutions for Service Level Agreements (SLA), Policies and Load Balancing Algorithms for Service Allocation, and Security.

Cloud Computing

On the other hand, Cloud layer does not aim to be bound to a specific provider, but instead provided in an interoperable fashion enabling freedom of choice depending on the specific needs. Therefore, this layer aims at making use of capacities which can be a combination of public, private or hybrid cloud solutions at IaaS and PaaS levels available in a wide diversity of providers and Cloud solutions.

- Cloud provides “unlimited” computing, processing and storage capacity however such unlimited potential is constrained by network accesses, since typically a latency of 80-100 micro sec is incurred to reach a cloud provider[i].
- Cloud platforms are particularly beneficial for computing intensive processing such as analysing large amount of data for instance
- Cloud platforms can also be used to provide and publish centralized services (APIs, Data-Lakes, Micro-Services) supported by the Asset Management, Massive storage and for coordination of the overall platforms.

Edge Computing

Edge Computing is the combination of digital capabilities which are connecting, integrating and interacting with physical devices to collect data and track events. Based upon specific application control, physical phenomenon behaviour is managed and influenced in almost real time (e.g. industrial internet, vehicles security, robots) using actuators rather than providing specific insights to end users aiming to take or influence decisions.

- The topology of an Edge system is complex and multiple Edge systems can be combined and inter-connected among them and to diverse clouds.
- The configuration, behaviour and technology behind Edge Computing are aimed at optimizing two main critical resources: QoS and latency, in particular.

The following subsections present in more detail building blocks and features envisaged at these three cooperating levels in the Federated Cloud Reference Architecture.

Edge Service Management

- The Edge Computing level is responsible for provision of computational and storage capabilities of the resources located at the Edge of the network. It is important to note that the locality of such deployment is typically bound to a specific space. This permits to offer extreme low latency that allows (near) real time data processing enabling data consumption and actuation in connected devices. Differently from existing commercial offerings where Edge installations rely on a single device, our architecture considers the possibility of establishing clusters of these Edge devices (which could be nested). The Edge layer is structured in three different levels:
- Connected Things represents IoT devices and Sensors that are the main data sources, but also are enabled with rapid actuation capacities and significant compute capacities in some cases;
- Things networking and Connectivity offer the connectivity services that permit extreme low latency.
- Edge Service Orchestration components enable elastic computing and storage capacities at the Edge, as well as, access to traditional Cloud computing capacities.

One of the main characteristics of Edge computing is the heterogeneity of its devices. Particularly the envisaged characteristics in Cloud Federation target the so-called large devices (using HEADS research project classification[25]) as in the following classification or Edge Servers:

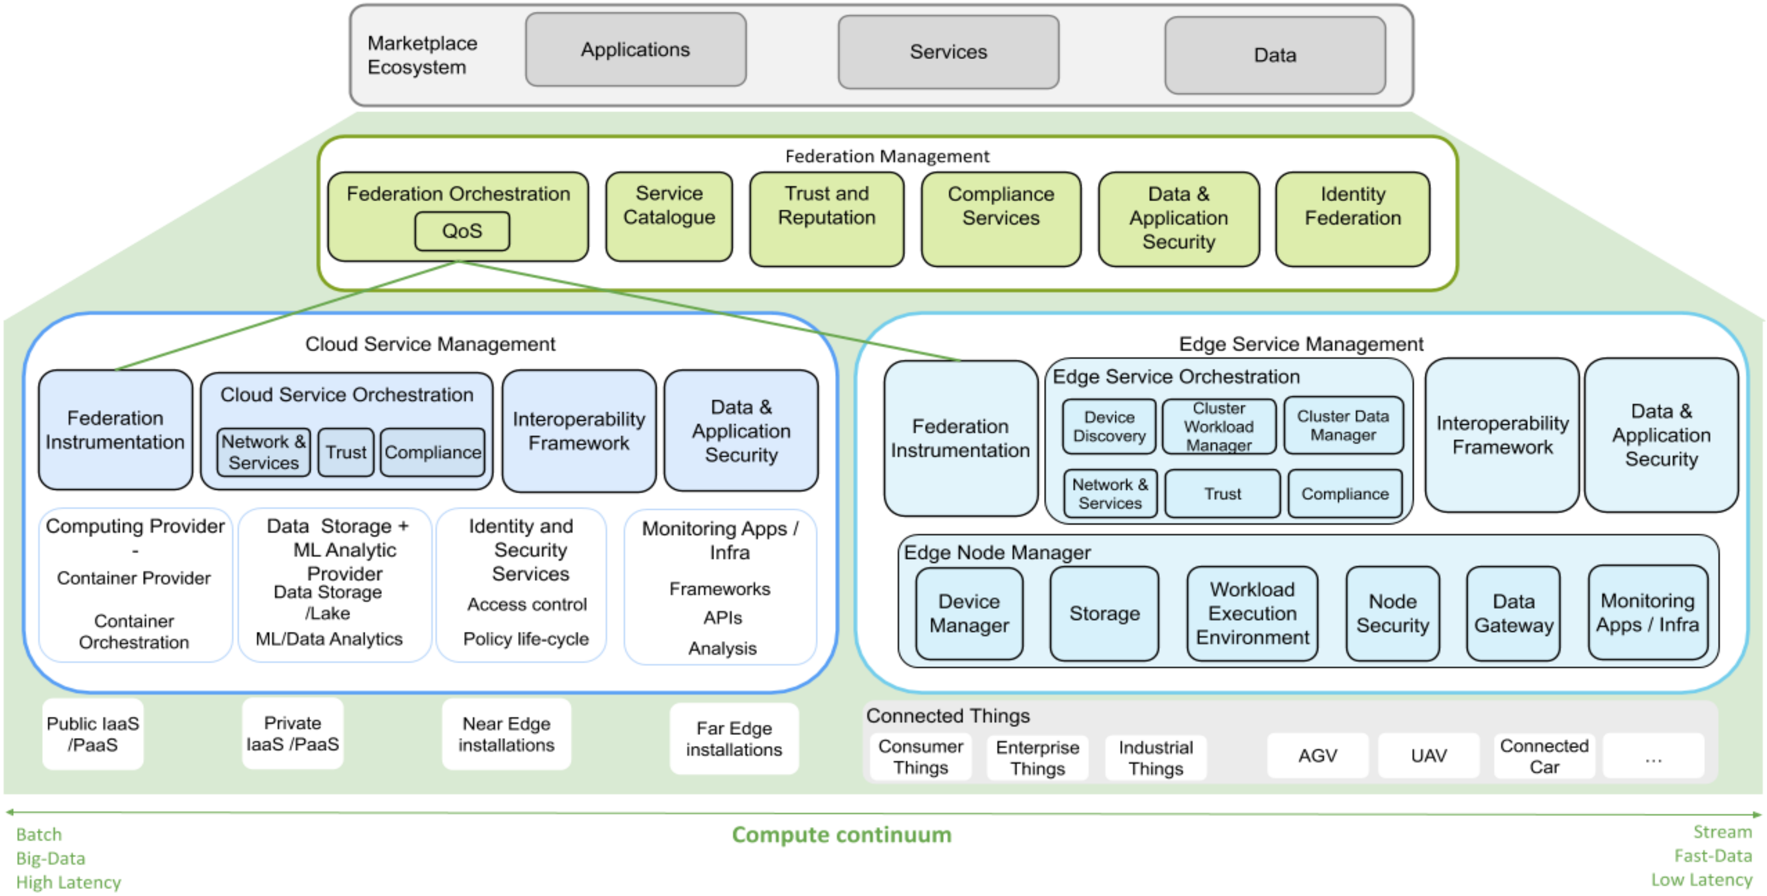

Figure 5 Cloud Federation Reference Architecture

- Tiny: Very limited devices (8 and 16 bit microcontrollers with less than 64kB program memory and 4kB of data memory). Example of this type of device is the Arduino UNO.
- Small: Devices with a specific OS and restricted hardware characteristics (less than 128kB program memory and less than 64kB data memory).
- Large: Devices supporting general purpose OS. Examples of these are: Raspberry PI and Android. Although still in its infancy, it is expected that under this classification we will soon find mini-HPC, bringing high performance computation at the edge of the network.

Current Edge computing environments consider Edge computing to be provided mainly from individual Edge servers deployed in the same locations that data sources. However, we anticipate that soon additional complexity levels at the Edge taking the consideration to form Clusters of Edge devices at the Edge. In order to have the ability to set-up Edge Clusters of large Edge devices,, we foresee the following components and functionalities:

- **Federation Instrumentation:** This component provides the means for Connected Things and Edge's computing and storage resources to participate in a service network. It manages Edge Cluster level operation of the established networks (including processes for monitoring, evaluation, runtime adaptation, etc.). In addition, this component delivers unified description for the Edge resources characteristics and capabilities, and their offered assets (i.e. resources, services and data), so to develop an abstraction layer for management of these resources. By means of such abstraction, the different interface implementations and characteristics of the diverse resources can be handled uniformly from the overall Federation Management.
- **Edge Service Orchestration: T**his set of components offer complete functionality for Edge Cluster management. The cluster needs to offer pre-configured solutions for diverse services to allow users to perform diverse operations with i.e. Data collected from Connected Things. A specific set of Services that can be executed in a Cloud Federation installation. In addition, the following features are envisioned for this block:
    - **Device Discovery:** This building block addresses communication patterns and protocols necessary to establish collaborative and resource discovery mechanisms. It enables the necessary mechanisms to make available to the Federated Cloud infrastructure Edge resources deployed at a specific location.
    - **Edge Cluster Workload manager** allows configuration of Edge cluster behaviour (controls, scheduling, access, fault tolerance and service life-cycle management) over services to be executed in the cluster. These could be packaged among others as containers or software functions. This component is responsible for deploying and managing the services in Edge infrastructures which can be composed of multiple resources. It enables allocating services available in the Service Registry to the most suitable resources producing optimal performance and efficient use of edge resources. Lifecycle Management (LM) will control the different execution phases of a Service. SLA Management will guarantee end users expectations regarding service performance and quality. The flow of actions within this module is the following: First, this module deploys these services in the most suitable resources available (discovered via Device discovery). and then it manages the lifecycle of services execution and termination of them. Finally, it

    also manages the QoS of the service being executed, being able to perform pre-defied corrective actions in case the expected QoS is not achieved.

  - **Edge Cluster Data Manager:** Relying on existing Object storage and NoSQL databases, these components offer different storage services to be used by Edge Node Storage systems to push and pull data, acting as an Event hub for all objects aligned to the specific Edge Cluster.
  - **Network & Services**: The purpose of this component is to offer tools and mechanisms which allow dynamic connection of devices at the Edge as well as to handle the deployment and scale secure virtual network services for these Edge devices.
  - **Trust**: This component permits to assess the trust levels with regards to the adherence to the established compliance frameworks as well as SLAs of both Edge devices and external Cloud providers, providing recommendation mechanisms to Edge Cluster Workload manager components.
  - **Compliance**: Edge is today recognised as an enabler for compliance, for instance, for GDPR, by enabling processing of data in its original generation location. However, diverse compliance frameworks can be of application in a Edge-cluster installation in highly regulated sectors such as eHealth, Energy or Banking. The role of this component is to monitor and enact that data processing is executed in compliant Edge devices.
- **Interoperability Framework:** This component enables having a single access point to interact with both computing and storage services in traditional cloud set-ups. It is worth to mention that it is expected that it offers cloud interoperability mechanisms, similar to the ones offered by jclouds[26] Apache Libcloud[27] or Libretto[28],for handling diverse cloud providers with the same interface so as to avoid lock-in to specific platforms of choice.
- **Data & Application Securit**y: Data & Application Security components at the Edge consider all appropriate mechanisms and tools in order to secure Edge data and application access and control. In addition, this can incorporate the use case specific necessary controls such as proactive threat detection technologies, vulnerability management and device and software patching cycles.
- **Edge Node Manager:** The Edge Node Manager component oversees the interface in order to manage a node, handling the resources at the Edge device. The main aim of this component is to deliver a unified description for the Edge resource characteristics and capabilities, and their offered capacity to the rest of the infrastructure. The Node manager develops an abstraction layer for all types of resources in the infrastructure which permits heterogeneous resources to be handled uniformly in the SAE instance.
  - **Device Manager:** Edge device manager is foreseen as a component which allows the automated software operation not specific for user services. These can include among others: update and rollout of baseline SOs, firmware, drivers and security features.
  - **Storage**: The common approach to persistent Edge storage is based on pre-existent distributed database systems which are now accommodated on the edge infrastructure. This approach however, needs to take into account specific requirements such as: data locality – low latency access time is a mandatory requirement; ability to scale beyond device capacity; mechanisms data synchronisation for across edge devices and cloud data

able to work off-line and resilient to network failures; need to strong integration with data node security.

- **Workload Execution Environment:** This component provides the means to perform actions related to the life-cycle of application components. The workload virtualization is based on containers, which facilitate the unified execution in heterogeneous execution environments in a variety of Edge devices. The most popular containerization system, Docker is now present in diverse constrained environments such as Raspbian Raspberry Pi Operating System, or Robot Operating System as well as specific network devices. Docker is increasingly being complemented with even more lightweight implementations of the containerization technologies among which are included Unikernels[29], Kata Container[30] and gVisor[31]. This permits us to predict the feasibility of our approach in even more constrained execution environments to those able to support Docker today.
- **Node Security:** Node Security component has to address the overall security challenges of Edge devices, data that is stored out of its operation as well as its access and transmission at Edge cluster level and as part of SAE instance operation to Cloud environments.
- **Data Gateway:** Offers an abstraction layer among diverse connected devices and the rest of interacting components. This component has to provide a unified interface that enables the collection of monitoring information about sensors' status, pulling and pushing data, interaction with actuation methods by a series of extensible plug-ins that covers the widest possible range of technologies and standards. Candidates for initial plug-in consideration are CoAP, MQTT and generic HTTP/HTTPS/Websockets protocols.
- **Monitoring:** It collects monitoring information about the status of the Edge node. The compiled parameters include the following aspects: physical infrastructure (memory, CPU usage and available storage) bandwidth (connection type and transmission rate), as well as, energy consumption. While bandwidth parameters provide a clear understanding of the quality offered by the node to the rest of the Cloud Federation, the energy consumption could be particularly important for certain scenarios that consider the optimisation of energy as a key aspect.

Cloud Service Management

The Cloud level aims to enable interoperability among diverse Cloud computing offerings while bringing highly performant resources for computation and data lakes. The components envisaged in this layer are the following:

- ***Federation Instrumentation*:** It provides the means for Multi-cloud orchestration platforms to abstract the services offered by the different Cloud provider's services. Similar to the equivalent functionality at the Edge Level it offers a common abstraction employed by Federation Management Capabilities. More in detail, Federation instrumentation at Cloud level has to offer a description framework for the diverse cloud services characteristics and capabilities, providing an abstraction layer for all the types of cloud services participating in the Federated Cloud platform. This will raise the need for defining a service description language which must be able to cover both functional and non-functional aspects of different cloud Services. This language should be able to describe them in different levels of abstraction and quality of service characteristics and infrastructure characteristics, such as levels of energy consumption.

- **Cloud Service Manager:** Cloud Service management tools control the execution of an application in Cloud service providers. These tools ensure that the agreed SLA is enforced and if not, takes the decision about the migration of the application components to a different infrastructure provider. They also use monitoring information to detect virtual infrastructure failures and apply defined elasticity rules to scale the virtual infrastructure assigned to application execution up and down. Service Registry at this level will present the preconfigured services and services templates can be available to be executed in the reference Cloud provider.
  - **Network & Services**: The Federated Network Manager is the software component that allows to build a federated cloud network by aggregating two or more Federated Network Segments, which are virtual networks within a cloud infrastructure, each sustained by a physical network backbone. The Federated Network Manager provides a uniform interface for users in order to set up a virtual federated network in a transparent way, independently from the underlying clouds. In order to do this, it features an API to allow for federated network definitions and uses adaptors to talk to the Cloud Management Platforms.
  - **Trust**: Trust as concept, is subjective in its nature, it is a multifaceted issue related to many other aspects such as risk, competency, security, perceptions, degree of utility and benefit, expertise, and previous experiences. In the context of a Cloud provider trust levels has to reflect how it performs and fulfils base requirements for the offered services. For building customer's trust, interesting criteria include the capacity to support the offered QoS and SLA levels together with the adherence to the defined compliance services. With this purpose, Trust services for Cloud providers need to offer mechanisms based on operational data that allow users transparency to validate these criteria.
  - **Compliance**: Compliance Services at the Cloud layer aims at ensuring to the user the Cloud providers obedience to certification schemes, standards and applicable codes of conducts. Among these SWIPO CoC, GDPR, EU Cloud Certification scheme could be considered contemporary examples, together with specific sector regulations. Similarly, to the analogous services at the Edge, the purpose of these services is to offer the Cloud user the mechanisms that allow to assess compliance with these applicable standards and schemes and regulations.

This component orchestrates diverse types of services existent in the Cloud provider:

- Computing Provider: Offering traditional IaaS and PaaS cloud capabilities.
- Data Storage and ML Analytic Provider: Offering both Block storage, Databases as a Service and Data Analytics capabilities including machine learning and pure analytics frameworks.
- Identity and Security Services: Both public cloud offerings and private cloud toolsets have in place mechanisms for identity management and security services to which Cloud Service managers need to interact.
- Monitoring: Monitoring services need to offer users of Federated Cloud services with the necessary information to ensure transparency over the status of execution of its services. This refers both to services that enable to gather real time with the status of employed infrastructures and services but also aspects such regulatory and security controls (incident inventory,

handling and corrective actions, disaster recovery plans, business continuity, etc.) and data location traceability.

- **Interoperability framework:** A significant barrier to Cloud computing adoption is interoperability and portability across providers and cloud services. In the scope of Federated Cloud Interoperability and portability of applications among Cloud environments require the ability to migrate applications across Cloud offerings. At the current state of affairs there are diverse challenges: Lack of adoption of standardised Cloud APIs; Diversity in services offered by Cloud providers languages, Non-interoperable accounting, billing, metering, and advertising services. Cloud providers in order to participate into Cloud federations, will need to agree on a set of services that will be offered to the Cloud Federation infrastructure and identify the standard definitions to provide. In addition, common formats for metered usage will need to be offered as part of the Interoperability framework.
- **Data & Application Protection Services:** Certain applications will require determined data and application protection to be built on top of services offered by Cloud providers other data protection services need to rely on built-in capacities of cloud providers. These application protection services can include encryption of data volumes attached to virtual infrastructure, services for data-base encryption and a service enabling object-level encryption for object-storage platforms.

Federation Management

Federation capabilities address the need for overall resource and service orchestration as the main aspect to be addressed in the Federated Cloud reference architecture. Main components for Federation End-to-End capabilities are:

- **Federation Orchestration** has to be provided as an open distributed runtime environment for decentralised management and coordination of multiple and diverse Edge devices and Cloud infrastructures. Federation management relies on Edge and Cloud Federation instrumentation components in order to implement at local level coordination actions. Federation management will provide a decentralized management framework that will base decisions on collaborative attributes and an entities' objectives lifecycle properties. Among these we consider trust-, administrative-, location-, relationships-, information-, assets-, contextual- and environmental- lifecycle properties. Cloud Federation will be based on the principle of subsidiarity of intelligence, meaning that decision making will happen at the lowest appropriate level, safeguarding autonomous behaviour of the participating environments.

  A key characteristic to be offered is service placement Optimization across the Edge to Cloud compute continuum: This best venue execution selection provides important benefits as the user is not locked into certain offers, but can decide the best choice for the deployment of each service among the different Edge and Cloud offers. The selection can be optimized based on the type of service to be deployed and different application execution requirements including QoS parameters and energy consumption execution requirements.

  A user may even want to launch a cross-site application that will be deployed in different venues at the same time. This can be beneficial for the distribution of the different components of a given application among different Edge and Cloud providers due to the application needs or across Clouds i.e. to improve fault tolerance in case of Cloud service disruption.

  Two of the main technological challenges of this approach are the creation and management of cross-site private networks and LANs using simple, standard procedures to interconnect different service components, and the creation and management of virtual storage systems across site boundaries to store service data.

Once Federation Orchestration services have taken the decision on best deployment venues for Service's components, Service's management processes take care of the deployment and operation observation of services in the different Edge and Cloud infrastructures. Within this context, service deployment tools will enable the packaging of services as well as the addition of security capabilities. Finally, this component needs to observe the execution of services by utilizing policies for harmonizing all management activities throughout the services lifecycle. These policies can integrate disparate service management requirements, from high-level Business Level Objectives to infrastructure requirements, into a unifying view to verify that the services are executing as expected. By focusing on service governance, it becomes easy to e.g., develop mechanisms for constant monitoring of service status and for triggering actions to increase and decrease capacity, i.e., enact elasticity rules, in order to meet the QoS specified for service execution.

  - Runtime adaptation needs to be incorporated to the Federation orchestration mechanisms by means of QoS management in order to trigger Federation evolution according to emerging situations, application requests and data protection risks, as part of context and environment characteristics.

- **Service Catalogue:** This building block has to be able to offer an integrated view of services offered by Cloud and Edge services available to the Federation.
- **Trust and Reputation:** Federated Cloud model depends on a high degree of the 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 Federated 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 a highly non-locked market of customers and providers where relationships are established on an on-off basis based on the mutual agreement. In addition, as the Federated 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 Federated Clouds services consumers' and providers' ecosystem.
- **Compliance services:** Federate Cloud scenarios are characterized by a constant flow of data which often cannot be specifically allocated to a particular geographical location. This results in uncertainty with regards to various data protection legislation, which may transcend national borders and therefore complicate compliance with global Data Protection legislation. Users of Federated Cloud Services can develop applications that handle confidential and private data. This stems from the need to safeguard its privacy. Therefore, from a legal point of view, providing mechanisms to ensure adherence to regulations and to enable data protection and privacy in Cloud environments should be a fundamental requirement. To ensure the ongoing integrity of this privacy, it is also important to contribute to and align with standards and policies created by industry organizations, commercial enterprises, and governments. In addition to generic data protection services, compliance services need to offer extensibility mechanisms in order to develop compliance and assurance models, assessment and recommendation support for sector specific regulatory frameworks. These need to offer particular emphasis on high-availability, data confidentiality, data/system/application/service integrity and auditability.
- **Data and Application Security:** While Federated Clouds offer a paradigm shifting technological solution for computational resources, data and application, the concerns about privacy and confidentiality of data as well as the underlying security and resilience of resources delivered in the Federated Cloud can prevent enterprises to widely uptake this delivery model. Federated Clouds

need to provide the necessary tools and mechanisms to handle security data across multi-vendor Edge and Cloud provider offers. These mechanisms need to include among others access control, configurable policies and mechanisms for auditability.

- **Identity Federation:** In Federated Cloud context it is crucial for participant Cloud and Edge providers to hold control over user's identities and credentials, creating the necessity for Federated Cloud environments to offer mechanisms for Federated identity management making use of existing standards that need to be agreed among participant entities.

4.2 Research Projects Positioning

Figure 6 details a high level view of the scope of contributions of Future Cloud Cluster research projects to the three levels of the Future Cloud Reference Architecture. Table 4 provides a more detailed mapping of the concrete building blocks in which each of the research projects has concrete results.

Afterwards a brief summary and links to each research project information are provided.

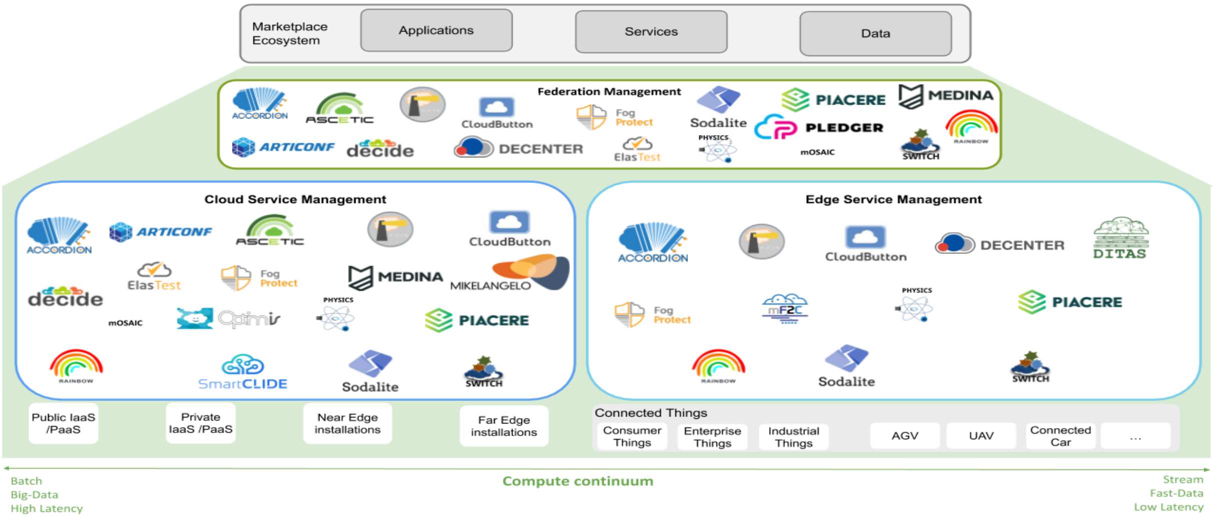

Figure 6 Mapping of Research projects to Reference Architecture layers

ACCORDION, https://www.accordion-project.eu/

Edge computing is transforming the utility computing paradigm: from centralized datacenter to highly distributed infrastructures. Edge computing is reducing the distance between user and server and minimising latency. Compared to a more traditional utility computing approach (i.e. based on clouds), edge computing is more 'democratic' as a consequence of its distributed and localised nature. It is quite trivial thus that many independent analysts estimated that edge computing will play a leading role in coming advanced technology. ACCORDION considers that by associating edge computing with advanced technologies such as 5G, the EU will be able to capitalise on its local resource and infrastructure and bring benefit to its SMEs. The project settles a practical approach in connecting, indexing and managing edge resources and infrastructures to support the so-called next-generation applications. Namely, those applications characterised by a set of requirements (e.g., latency, bandwidth, data localization) that make unfeasible their execution on remote clouds. To this end, ACCORDION establishes an opportunistic approach in bringing together in a highly heterogeneous federation edge resources and infrastructures (public clouds, on-premise infrastructures, telco resources, end-devices) that can support NextGen application requirements. ACCORDION orchestrates the compute & network resources that organize in a federative continuum spanning from edge to public clouds.

Table 4 Mapping of Research projects to Reference Architecture Building blocks

			ACCORDION	ARTICONF	ASCETiC	BEACON	CloudButton	DECIDE	DECENTER	DITAS	Elastest	FogProtect	MEDINA
Federation Management	Federation Orchestration		x	x	x	x	x		x				
	Service Catalogue				x			x					
	Trust & Reputation			x				x					
	Compliance Services							x			x	x	x
	Data & Application Security			x		x						x	x
	Identity Federation			x									
Cloud Service Management	Federation Instrumentation		x	x			x	x			x		
	Cloud Service Orchestration	Network & Services	x	x	x	x		x			x		
		Trust		x	x								
		Compliance			x						x	x	x
	Interoperability Framework		x		x	x	x						
	Data & Application Security		x								x	x	x
Edge Service Management	Federation Instrumentation		x						x	x			
	Edge Service Orchestration	Device Discovery											
		Cluster Workload Mng.	x				x			x			
		Cluster Data Mng.					x			x			
		Network & Services	x			x							
		Trust											
		Compliance								x		x	
	Interoperability Framework						x						
	Data & Application Security											x	
	Edge Node Manager	Device Manager	x							x			
		Storage	x							x			
		Workload Execution Env.							x	x			
		Node Security							x	x		x	
		Data Gateway								x		x	
		Monitoring	x						x	x			

			mF2C	MIKELANGELO	mOSAIC	OPTIMIS	PHYSICS	PIACERE	PLEDGER	RAINBOW	SmartCLIDE	SODALITE	SWITCH
Federation Management	Federation Orchestration				x		x		x			x	x
	Service Catalogue				x			x		x	x		
	Trust & Reputation							x					
	Compliance Services				x								
	Data & Application Security										x		
	Identity Federation												
Cloud Service Management	Federation Instrumentation					x					x	x	
	Cloud Service Orchestration	Network & Services		x	x		x	x		x	x	x	x
		Trust		x		x					x	x	x
		Compliance			x	x					x		
	Interoperability Framework			x	x	x	x					x	x
	Data & Application Security				x	x		x	x		x	x	
Edge Service Management	Federation Instrumentation		x						x	x		x	
	Edge Service Orchestration	Device Discovery	x					x		x		x	
		Cluster Workload Mng.	x				x			x		x	
		Cluster Data Mng.								x			
		Network & Services					x			x		x	
		Trust							x	x		x	
		Compliance											
	Interoperability Framework		x				x						
	Data & Application Security									x			
	Edge Node Manager	Device Manager	x							x			
		Storage								x			
		Workload Execution Env.	x				x		x	x			
		Node Security	x							x			
		Data Gateway								x			
		Monitoring	x				x	x	x	x		x	

To this end, ACCORDION establishes an opportunistic approach in bringing together in a highly heterogeneous federation edge resources and infrastructures (public clouds, on-premise infrastructures, telco resources, end-devices) that can support NextGen application requirements. ACCORDION orchestrates the compute & network resources that organize in a federative continuum spanning from edge to public clouds. Deployment decisions are based on a multi-criteria analysis based on AI-based approaches and are taken by ensuring that requirements on privacy, security, cost, time and resource type are matched. To support deployment decisions, ACCORDION services include specific functionalities devoted to resource indexing, monitoring and organization. Edge resources are grouped and organized using a specifically derived abstraction named Minicloud. The minicloud itself defines and provides services for the actual management of applications. ACCORDION monitors applications running at the edge and in cloud assessing the QoE provided to end-users. When performances degrade and QoE is affected, ACCORDION takes decisions triggering the reconfiguration of applications and their allocation plans on resources. ACCORDION also provides specialized services for edge and cloud storage management.

ARTICONF, https://articonf.eu/

Co-located and Orchestrated Network Fabric (CONF) is an automated cloud virtual infrastructure solution for social network applications developed in the EU H2020 ARTICONF project. The CONF framework extends the software developed in the SWITCH project, and aims to improve the existing infrastructure support in the DevOps lifecycle of social network applications to optimize QoS performance metrics as well as ensure fast recovery in the presence of faults or performance drops. The CONF aims to complement existing IaaS (Infrastructure-as-a-Service) solutions with flexible infrastructure planning and provisioning solutions and provide rich support for the high service quality and trustworthiness required by social network applications.

Besides infrastructure automation solutions[32] [ariconf-1], the CONF framework improves trustworthiness of the SLA during the service operation using blockchain and smart contract technologies.The basic idea is shown in Fig. articonf-conf. When the SLA is initialised (step 1), two smart contracts are also created for automating the SLA (step 2a), and for monitoring (step 2b). The user accepts the SLA but also is willing to pay the fee for monitoring (step 3a and 3b). We developed a game theory-based incentive model to attract volunteer witnesses for joining the monitoring of the SLA (step 4a), who report the violations based on their monitoring results (step 4b). In this way, the SLA will be enhanced by the real-time monitoring and violation report done by the witness (step 5). The detailed algorithm have been published in IEEE INFOCOM 2019[33].

ASCETiC, http://ascetic-project.eu/

ASCETiC project focused on issues of energy efficient computing, specifically on design, construction, deployment and operation of Cloud services. It proposed 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.

ASCETiC goal was to characterize the factors which affect energy efficiency in software development, deployment and operation. The 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 measurements of energy requirements into the design and development process for software to be executed on a Cloud platform.

ASCETIC main advantage is to provide Cloud customers and providers with the ability to determine and optimize the relationship between energy consumed by an application, and its delivered gain. This enables to determine gained value per each unit of energy spent in IT, and to correlate energy consumed to the beneficial work for which that energy is spent.

ASCETiC Toolbox functionalities consider three layers: SaaS layer facilitates the 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; IaaS layer considers the admission, allocation and management of virtual resources.

BEACON, https://cordis.europa.eu/project/id/644048

With the increasing amount of infrastructure cloud services becoming available there are many benefits to interconnecting several cloud services in the form of a federation. There is a strong industry demand for automated solutions to federate cloud network resources, and to derive the integrated management cloud layer that enables an efficient and secure deployment of resources and services independent of their location across distributed infrastructures.

The BEACON framework for federated cloud networking is based on an associated reference architecture for the federated cloud networking. BEACON provides many advanced features for the federated cloud networks such as automated high availability, datacenter location-aware elasticity, automated Service Function Chaining (SFC), and security across clouds.

The proposed federation network model addresses the challenge of federating clouds based on different cloud management platforms and network technologies (i.e. Software Defined Networking, SDN)2 . Moreover, it can be used in different cloud federation architectures such as peer, hybrid, and brokered cloud federation.

The BEACON Network Manager is the software component that allows to build a federated cloud network by aggregating two or more Federated Network Segments, which are virtual networks within a cloud infrastructure, each sustained by a physical network backbone. The BEACON Network Manager provides a uniform interface for users in order to set up a virtual federated network in a transparent way, independently from the underlying clouds. In order to do this, it features an API to allow for federated network definitions, and uses adaptors to talk to the Cloud Management Platforms. The BEACON Network Agent drives the control plane of a federated network. It informs other Network Agents about the known network segments of its domain, and instructs the BEACON Datapath. The BEACON Datapath defines the data plane as instructed by the Network Agent. The Datapath encapsulates traffic between the different Federated Network Segments and provides the needed mapping to interconnect Federated Network Segments.

CLOUDBUTTON, https://cloudbutton.eu/

This project is inspired by the following sentence from a professor of computer graphics at UC Berkeley : “Why is there no cloud button?” He outlined how his students simply wish they could easily “push a button” and have their code – existing, optimized, single-machine code –running on the cloud.”

The main goal is to create CloudButton: a Serverless Data Analytics Platform. CloudButton will “democratize big data” by simplifying the overall life cycle and programming model thanks to serverless technologies. To demonstrate the impact of the project, we target two settings with large data volumes: bioinformatics (genomics, metabolomics) and geospatial data (LiDAR, satellital).

While serverless computing is changing the way, applications run in the cloud, vendor lock-in limits the adoption of multiple serverless computing platforms. Furthermore, the complexity of different Cloud APIs preclude massive adoption of Cloud platforms. In this project, we advocate for access transparency: enabling local and remote resources to be accessed using identical operations. To this end, CloudButton presents Lithops: a novel Multi-Cloud toolkit that allows to run regular, multiprocess Python programs at scale. Our toolkit completely hides the underlying complexities of the different serverless computing providers, enabling access transparency for end users to the multicloud. CloudButton demonstrates how the toolkit is a novel building block capable of constructing different types of multiprocess-based applications, and execute them using the major serverless computing vendors in today's market such as Amazon, Google, and IBM.

CloudButton contributes to the federation of Clouds with its multi-Cloud Architecture, where our Lithops toolkit supports different Backends for Compute and Storage resources in many Cloud providers. The concept of access transparency, and mapping local resources (processes, storage) to remote ones is essential in a multi-cloud setting. We can easily invoke and combine heterogeneous resources from different backends using Lithops. Furthermore, CloudButton has actively contributed to the orchestration of heterogeneous resources thanks to trigger-based orchestration of serverless workflows (TriggerFlow). We plan to contribute to federated Cloud orchestration mechanisms, cloud continuum optimizations, and programming language toolkits enabling access transparency for different applications.

DECENTER, https://www.decenter-project.eu/

DECENTER designs and develop a robust and secure Fog Computing Platform, based on existing Open Source frameworks that provides application-aware orchestration of resources and zero-touch provisioning of Microservices for Artificial Intelligence applications based on their specific requirements and described by means of Quality of Service models.DECENTER relies on a framework that catalyses offer and demand for Edge resources. The innovative solution is based on Blockchain to draft and seal Smart Contracts, monitor their fulfillment and grant rewards.

DECENTER evolves decentralised Artificial Intelligence models that become possible with a more reactive and flexible resources infrastructure, unlocking a huge potential for innovative applications. Among others, it addresses the needs for data locality and time-critical aspects of the Big Data problem.

DECENTER offers support for federation of edge/Cloud resources across multiple administrative regions. It improves the resource advertising and leasing through the network.In addition it offers deployment and monitoring of a monitoring of SLAs ensuring trust in the system.

DECIDE, https://www.decide-h2020.eu/

DECIDE created a DevOps framework based on the concept "Extended DevOps approach", and focusing not only on the basic characteristics of DevOps, namely CI and CQ, but also in the architectural phase, pre-deployment and operation. The applications targeted were multi-cloud applications whose components are deployed on multiple CSPs and with strict non-functional requirements such as availability, performance, cost and legal aspects. Among the different outcomes that have resulted from DECIDE, there are three that can be of interest for a federation of clouds. These are:

- ACSmI (Advanced Cloud Service metaIntermediator): ACSmI Discovery is a catalogue of cloud services that DevOps teams can use to choose the most adequate ones for them. The discovery can be performed through a simple search / filter but also through a multi-objective multi-criteria optimization tool (OPTIMUS). ACSmI Monitoring monitors the instances deployed on the service offerings and compares them with certain SLOs of their SLAs, as agreed. Whenever a SLO is not fulfilled, an alert is raised.
- ADAPT: This tool allows to automatically deploy and redeploy the components of the application on multiple cloud services. To automatically deploy, the machines are provisioned, policies applied and all necessary software installed, such as the containers of the application or agents for the monitoring of the performance VMs. ADAPT also monitors the components of the application through a safe-method approach and assesses whether the application agrees with the application defined SLA or not, and in the latter case, an alert is raised for an (automatic) redeployment to be performed..

– OPTIMUS: multi-criteria multi-objective optimization tool to deploy a microservices application on multiple CSPs

ACSmI can be placed under the Federation Instrumentation, federation orchestration, service catalogue and service management in the reference architecture above.

DITAS, https://www.ditas-project.eu/

The DITAS Cloud Platform allows developers to design data-intensive applications, deploy them on a mixed cloud/edge environment and execute the resulting distributed application in an optimal way by exploiting the data and computation movement strategies across Edge and Clouds. Data has to be moved on demand between edge and cloud, partial and updated data also needs to be kept in sync and simultaneously the data has to go through different transformations depending on the user requirements. The concept of Virtual Data Containers, which let the developers simply define the requirements on the needed data, expressed as data utility, and take the responsibility of providing these data timely, securely, and accurately by hiding the complex underlying infrastructure composed by different platforms, storage systems, and network capabilities.

To achieve this goal, Virtual Data Container implement data and computation movement strategies to decide where, when, and how to save data – on the cloud or on the edge of the network – and where, when, and how to compute part of the tasks composing the application to create a synergy between traditional and cloud approaches able to find a good balance between reliability, security, sustainability, and cost.

DITAS provides a dedicated SDK for data administrators, application developers and architects and operators that simplifies the life of the actors involved in the management of the related VDC according to the process described. Depending on the role, the SDK is provided in different flavors: e.g., CLI, GUI, web applications.

Elastest, https://elastest.io/project.html

ElasTest is a platform aimed to ease end to end testing of distributed systems. The two key features of the platform are: 1) Provide an easy deployment process and easy access to the necessary services usually involved in an end to end test and 2) Provide easy-to-use tools to show and analyze logs and metrics of all elements involved in an end to end test.

New advances in ICT technology influence the way software is developed and tested, the proliferation of large-scale applications targeting thousands of users that can be connected concurrently and expect real time interactions; makes the testing strategy a crucial aspect for the release management process of the applications. Nowadays cloud technologies are creating advantages for organizations that adopt it such as: speed, agility, scalability, accessibility and flexibility; therefore ElasTest aims to extend the adoption of the aforementioned benefits offered by the cloud to testers through the creation of a cloud platform (ElasTest Platform) designed for helping to validate large software systems that require complex test suites and validation processes.

The ElasTest cloud components are concerned with the management and monitoring of the resources that the platform needs to operate; as well as of the lifecycle management associated to the on-demand testing support services catalogue which can be requested by the ElasTest Platform user dynamically. In addition to the cloud components in charge of the platform management, the report also includes other kind of cloud-based component not targeting the platform itself but offering management capabilities over the software system under evaluation.

FogProtect, https://fogprotect.eu/

Cloud computing is transitioning from a few large data centres to a truly decentralized computing paradigm. Computing resources are increasingly provided near the network edge, in a geographically distributed way, in the form of so-called fog nodes. Data produced in end devices like smartphones, sensors or IoT devices can be stored, processed and analysed across a continuum of compute resources, from end devices via fog nodes to cloud services. This decentralized computing paradigm provides huge benefits in terms of reduced latency, increased processing speed and energy savings, but the protection of sensitive data in such a widely decentralized setting becomes a critical concern.

FogProtect delivers new and advanced architectures, technologies, and methodologies for ensuring end-to-end data protection across the computing continuum, from cloud data centres through fog nodes to end devices. The FogProtect solutions are generic and can be used in multiple contexts to support many types of applications and services. FogProtect combines four main technology innovations: (1) secure data container technology for data portability and mobility, (2) data-protection-aware adaptive service and resource management, (3) advanced data protection policy management, (4) dynamic data protection risk management models and tools.

FogProtect supports federation across the computing continuum, from cloud data centres through fog nodes to end devices. FogProtect manages data protection across the whole computing continuum, also reconciling it with other concerns like performance and energy consumption.

MEDINA, https://cordis.europa.eu/project/id/952633

Despite the evident benefits of cloud computing, its adoption is still limited partially because of EU customers' perceived lack of security and transparency in this technology. Cloud service providers (CSPs) usually rely on security certifications as a means to improve transparency and trustworthiness, however European CSPs still face multiple challenges for certifying their services (e.g., fragmentation in the certification market, and lack of mutual recognition). In this context, the new EU Cybersecurity Act (EU CSA) proposes improving customer's trust in the European ICT market through a European certification scheme.

The proposed EU CSA's cloud security certification scheme conveys new technological challenges due to its notion of "levels of assurance" (e.g., high-assurance through continuous certification for the whole supply chain), which need to be solved in order to bring all of EU CAS's expected benefits to EU cloud providers and customers.

In this context, MEDINA proposes a framework for achieving a continuous audit-based certification for CSPs based on the EU CSA's scheme for cloud security certification. MEDINA will tackle challenges in areas like security validation/testing, machine-readable certification language, cloud security performance, and audit evidence management. The MEDINA consortium is composed of academic and industrial partners, which play key roles in the EU cloud security certification ecosystem (e.g., research, cloud providers/customers, and auditors). MEDINA will provide and empirically validate sustainable outcomes in order to benefit EU adopters.

For the federation of clouds, MEDINA provides a solution to continuously monitor the compliance of the service with respect to the European cloud security certification scheme.

mF2C, https://www.mf2c-project.eu/

The main objective of the mF2C project was to design and develop a hierarchical, open, secure, decentralized and coordinated management platform facilitating the efficient usage of resources, taking into consideration service requirements and user demands, in a paradigm shifting scenario combining cloud and fog computing.

The F2C collaborative and coordinated computing ecosystem has been conceived to: i) efficiently and transparently utilize available distributed and heterogeneous resources at the edge; ii) support applications and services that do not fit well into the paradigm of the traditional centralized cloud (e.g., low latency, fast handover and connectivity of mobile applications), and; iii) pave the way to new business models in both cloud and smart devices sectors. Security and privacy are also addressed in a complementary fashion in F2C systems. The highly complementary features of security and privacy, throughout both cloud and fog, make the F2C system a highly interesting subject of future research and innovation. Last but not least, despite fog's potential, many devices are likely to suffer from resource issues (limited storage, battery, compute power), which can lead to inability to satisfy service level agreements. Hence, a critical question here is how can collaborative scenarios, based on resource sharing and clustering, extend the concept of cloud provider to an unknown frontier.

mF2C project has developed a the mF2C framework provides a coordinated management of traditional cloud architectures and novel edge ones, offering unique capabilities for distributed execution of applications throughout IoT, fog and cloud environments. mF2C can support the necessary federation across Edge and Cloud resources while considering complex Edge scenarios including aggregations of Edge devices and hierarchies of Edge clusters.

MIKELANGELO, https://www.mikelangelo-project.eu/

The Cloud and HPC architectures are a trade-off between efficiency, stability and security. Legacy and compatibility requirements have through the years amassed a significant amount of layers, on top of which the actual application code runs. This complexity requires complex setup and management tools that have to follow the development of many pieces of code. Having so diverse and large infrastructure results in large security attack surface, reduced only with sophisticated networking security measures.

MIKELANGELO provides improved responsiveness, flexibility and security of virtual infrastructure through the use of unikernels - OSv (osv.io). Using the unikernels, we target a unified architecture, with components supporting different variants of the HPC, Cloud and HPC-Cloud. MIKELANGELO relies on optimisation of guest, hypervisor and their joint collaboration. Unikernel OSv is used to reduce the size of guest virtual machine to a bare minimum, to reduce its complexity while offering the best support for legacy applications among the unikernels.

The technologies developed and offered within MIKELANGELO range from the improvements of the hypervisor, where introduced IO workload monitor (IOcm) facilitating dynamic optimisation of CPU core allocation for IO intensive workloads, a novel para virtualized IO model for KVM. On the security side, we developed a method dubbed Side-Channel Attack Mitigation (SCAM), preventing side-channel attacks on KVM hypervisor.

Finally, to improve the practicalities of using the unikernels - we developed a toolset to simplify the configuration and application management within OSv has been provided - enabling their simplified configuration and deployment.

mOSAIC, https://www.mosaic-cloud.eu/ , https://cordis.europa.eu/project/id/256910

mOSAIC - Open-Source API and Platform for Multiple Clouds - was motivated by low availability of programming models for Cloud applications, tools for easy deployment of applications in multiple Clouds, and user-driven service level agreements, as well as platform dependability and non-portability of Cloud applications ("cloud vendor lock-in") due to different APIs for different types of Cloud services.

mOSAIC aimed to: (a) design a language- and platform-agnostic application programming interface for using multi-Cloud resources; (b) build an open-source, portable and vendor-agnostic platform for using Cloud services based on the proposed API; (c) define a machine-readable Cloud Ontology and perform

Semantic based discovery of agnostic Cloud services and resources driven by Application and Cloud Patterns (d) build user-centric service level agreement, and resources and services brokerage, negotiation, monitoring and dynamic reconfiguration based on multi-agent technologies and semantic data processing; (e) build proof-of-concept applications.

The key results that distinguished mOSAIC from other solutions for multi-Clouds portability and interoperability are the following:

A new level of abstractions of the Cloud resources that allows not only an uniform access to multiple Clouds, but also to decouple from the inherited style of programming of the accessed services; the conceptual API is implemented currently in Java, Python and Erlang.

- Targeting the application developer, an entire set of tools was built for an easy design of the Cloud applications. Eclipse plug-ins, work-benches, templates, and various front-ends, like web interfaces are able to assist the developer. In particular the Semantic Engine and Dynamic Semantic Discovery Service support the user in discovering the resources and services offered by mOSAIC and various 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. Another software prototype is the Personal Testbed Cluster that allows the development, testing and debugging of the codes on own desktop, and then, with the help of the other tools, to experience a seamless deployment in multiple Clouds.
- The selection of the Cloud service to be consumed 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 connected;
- An open-source and deployable Platform-as-a-Service that 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, not encountered elsewhere.
- A set of innovative applications relying upon infrastructure and software services from multiple Clouds.

OPTIMIS, https://cordis.europa.eu/project/id/257115

OPTIMIS aimed at optimizing IaaS cloud services by producing an architectural framework and a development toolkit. The optimization covered the full cloud service lifecycle (service construction, cloud deployment and operation). OPTIMIS gives cloud service providers the capability to easily orchestrate cloud services from scratch, run legacy apps on the cloud and make intelligent deployment decisions based on their preference regarding trust, risk, eco-efficiency and cost (TREC). It supports end-to-end security and compliance with data protection and green legislation. It also gives service providers the choice of developing once and deploying services across all types of cloud environments – private, hybrid, federated or multi-clouds.

Fundamentally, the OPTIMIS toolkit is a cloud enabling technology that helps end-users create a truly digital IT infrastructure, by adding an automation and orchestration layer to a virtualized infrastructure. The OPTIMIS components are deployed in the data center and are a complement to cloud management and orchestration platforms. OPTIMIS enables users to schedule and automate the delivery of workloads to the most suitable clouds (internal/external) based on policies such as trust, risk, eco-efficiency and cost (TREC), as ‘best execution venue’ strategy.

OPTIMIS simplifies the management of infrastructures by automating most processes while retaining control over the decision-making. The various management features of the OPTIMIS toolkit make infrastructures adaptable, reliable and scalable. These, altogether, lead to an efficient and optimized use of resources.

By using the OPTIMIS toolkit, organizations can easily provision on multi-cloud and federated cloud infrastructures and allows them to optimize the use of resources from multiple providers in a transparent, interoperable, and architecture-independent fashion.

PHYSICS, https://cordis.europa.eu/project/id/101017047

PHYSICS empowers European CSPs exploit the most modern, scalable and cost-effective cloud model (FaaS), operated across multiple service and hardware types, provider locations, edge, and multi-cloud resources. To this end, it applies a unified continuum approach, including functional and operational management across sites and service stacks, performance through the relativity of space (location of execution) and time (of execution), enhanced by semantics of application components and services. PHYSICS applies this scope via a vertical solution consisting of a:

-Cloud Design Environment, enabling design of visual workflows of applications, exploiting provided and generalized Cloud design patterns functionalities with existing application components, easily integrated and used with FaaS implementations, including incorporation of application-level control logic and adaptation to the FaaS model.

-Optimized Platform Level FaaS Service, enabling CSPs to acquire a cross-site FaaS platform middleware including multi-constraint deployment optimization, runtime orchestration and reconfiguration capabilities, optimizing FaaS application placement and execution as well as state handling in collaboration with stateless functions, while cooperating with provider-local policies

-Backend Provider Optimization Toolkit, enabling CSPs to enhance their baseline resources performance, tackling issues such as cold-start problems, multitenant interference and data locality through automated and multi-purpose techniques.

PHYSICS will produce a Reusable Artefacts Marketplace (RAMP), in which internal and external entities (developers, researchers etc) will be able to contribute fine-grained reusable artifacts (functions, flows, controllers etc).

With relation to the proposed RA, PHYSICS may contribute in aspects that have to do with the Application Marketplace, through the RAMP implementation and the contributed artefacts. Furthermore, it creates mechanisms in the scope of the Federation Orchestration, including QoS monitoring and evaluation of available or candidate resources and services, middleware layer for orchestrating intermediate platform offerings (FaaS in particular) across multi-cloud and edge IaaS and container services , decision support (function and data) placement algorithms as well as intelligent cluster workload manager schemes for optimal local edge and cloud execution.

PIACERE, https://cordis.europa.eu/project/id/101000162

The growing role of software in managing infrastructures and the DevOps movement, focused on the automation of infrastructure management, are targeting the challenges of increasing speed and quality of infrastructure management, thus lowering costs and enhancing security and trustworthiness. However, the market of infrastructure automation tools is fragmented, there is no single one to manage the whole lifecycle of infrastructure as code (IaC) and existing solutions do not address all trustworthiness and security aspects throughout the whole lifecycle. PIACERE will develop tools, techniques and methods enabling organisations to fully embrace the IaC approach through the DevSecOps philosophy. PIACERE will provide the first Integrated Development Environment (IDE) to develop and verify IaC. Exploiting Model-Driven Engineering (MDE), the IDE will enable developers to create infrastructural code at an abstract level. Using the novel DevOps Modelling Language (DOML), the DevOps team will generate IaC for different languages and verify its correctness at model and code level along with the corresponding security components. The IDE is one part of the complete workflow and will be supported with: 1) a canary environment to aid the simulation of the conditions of the production environment allowing the early identification of potential vulnerabilities and 2) an IaC execution Environment to automatically deploy, monitor and ensure that the conditions are met, incorporating self-healing and self-learning features. The integration, security first and IaC polyglotism arm the DevSecOps teams to treat and work with IaC as they do with traditional code, simplifying the design, development and operation of IaC, while increasing their productivity, quality and reliability

For the federation of clouds, PIACERE can contribute in the federation orchestration, service management and management of the operational deployment of the cloud continuum application, with special focus on infrastructure.

PLEDGER, http://www.pledger-project.eu/

Current approaches on edge computing are not enough to address the forthcoming massive usage of edge computing in the frame of large Internet of Things (IoT) deployments. Essentially, the main goal in such scenarios is to ensure that the overall offered Quality of Service (QoS) fits the application needs over the edge or edge/cloud continuum deployments.

Processing speed and latency issues have been identified as the top barrier in this domain, while cost and reliability (meeting the provider Service Level Agreements - SLAs) are the top and second most important factors for evaluating edge and cloud services.

At the same time, network capabilities also play a very important role. Especially performance may hinder the adaptation of edge solutions for a variety of applications (e.g. in the Virtual/Mixed Reality applications), for which timely completion is key for the productivity of the company.

With this regards our project Pledger aims to research these topics and provide advancements using: Rich data sets for better Machine Learning, Consolidating data in the cloud for sophisticated Machine Learning applications; High Performance Computing, Leverage the Edge for computationally intensive analysis or machine learning; Security, Trusted and secured end-to-end communications; and Blockchain, Bringing Smart Contracts to the Edge.

RAINBOW, https://rainbow-h2020.eu/

The broad vision of the RAINBOW project is to empower IoT service operators to solely focus on the design and development of their services business logic, leaving to RAINBOW the burden of how and where services must be placed (in the fog continuum), establishing secure collaboration among entities and dealing with low-level aspects in data analysis including heterogeneous resource management, mobility and data movement.

Specifically, RAINBOW abstracts and seamlessly handles:

- The deployment, placement and runtime adaptation of geo-distributed IoT services through novel AI-driven orchestration algorithms capable of providing in-time and decentralized decision making for federated resource and network fabric administration to ensure user-defined QoS, energy and network optimization policies are met at all times.
- The establishment of “trust” among collaborating entities, while also verifying security primitives across the device-fog-cloud-application stack through a “zero-conf” overlay mesh networking

paradigm ensuring encrypted and dynamic package routing even through geo-distributed administration domains.

- Interoperable data processing across the fog continuum by pushing “intelligence” to the network “edge” with -in place- data management and fog analytics services through decentralized edge API’s capable of “talking and sharing” to each other without offline, manual and human intervention

SODALITE, https://www.sodalite.eu/

SODALITE is addressing heterogeneous, software-defined, high-performance computing environments by considering environments that comprise accelerators/GPUs from the software-defined point of view. SODALITE provides an IDE, enabling the developers and infrastructure providers with the less complicated tools to manage the aforementioned software-defined, application-specific infrastructures. They are able to declare their application and infrastructure requirements through a simpler descriptive language, while being supported by an intelligent semantic suggestion system and further extensively checked for possible errors, with an intelligent, pattern-based system. After the application and infrastructure have been successfully described, the application is optimised (based on the available source code and the targeted infrastructure), thus improved even before deployment. During application execution, an innovative application predictive refactoring system, coupled with monitoring, improves the application execution under certain conditions (e.g., if the application performance is degraded; better resources are available; pre-set constraints are not met).

Having described the SODALITE project, its clear positioning is as follows. It uses standardised TOSCA as the central, common language between components and as the blueprint for the application orchestration and infrastructure provisioning. It recognises the inherent heterogeneity of infrastructures and provides modelling support for them (even through automatic discovery of capabilities). It builds on the well-known and well-used principles such as ontologies for semantic support, code smell detection and provides the code as open source.

SODALITE, as it focuses on the holistic view on the application and its deployment on the infrastructure, can contribute in the federation orchestration and then in the orchestration and optimisation (at runtime) within Cloud and Edge infrastructures.

SWITCH, https://cordis.europa.eu/project/id/643963

The SWITCH framework tackles the challenges in developing and running time critical applications in cloud environments. The project focuses on three use cases: collaborative real-time business communication, elastic disaster early warning, and cloud studio for directing and broadcasting live events. Those application quality constraints, such as jitter and latency in live event broadcasting or processing delay for sensor data in disaster early warning, require not only advanced infrastructure, but also sophisticated optimization mechanisms for developing and integrating system components. The SWITCH workbench couples the development, deployment and operation phases of the time critical application lifecycle within cloud environments, aims to reduce the development complexity of building such applications and improve the efficiency of runtime application control in the application lifecycle via three key overlapped cycles of *application development*, *provisioning* and *runtime* via three subsystems[34],[35].

The SWITCH Interactive Development Environment (SIDE) subsystem provides interfaces for all of the user- and programmer-facing tools, by exposing a collection of graphical interfaces and APIs that tie in 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) identifying the constraints on infrastructure resources required to meet the time-critical requirements of the applications; 2) defining an optimal virtual runtime environment that meets those constraints; 3) provisioning the planned environment with the chosen resource provider; and 4) deploying the components required by the application.

The Autonomous System Adaptation Platform (ASAP): 1) monitors the status of the application and the runtime environment; 2) examines the actual performance of the required quality attributes; 3) autonomously controls the application and runtime environment to maintain optimal system level performance against the time critical constraints; and 4) learns from its own decision history to improve its intelligence in making future decisions for autonomous control.

4.3 Mapping of Future Cloud Research Areas to Architecture Building blocks

In August of 2020 in the Future Cloud Cluster we produced a revision of the research areas and challenges in the scope of Edge computing, Multi-Cloud, Computing continuum and Federated Cloud, as part of our H2020 Future Cloud Research Roadmap. The research areas and challenges express the vision of the cluster participants for future developments for Future Cloud.

We have articulated our views for future research in this context around thirteen research areas which aggregate 36 Research Challenges:

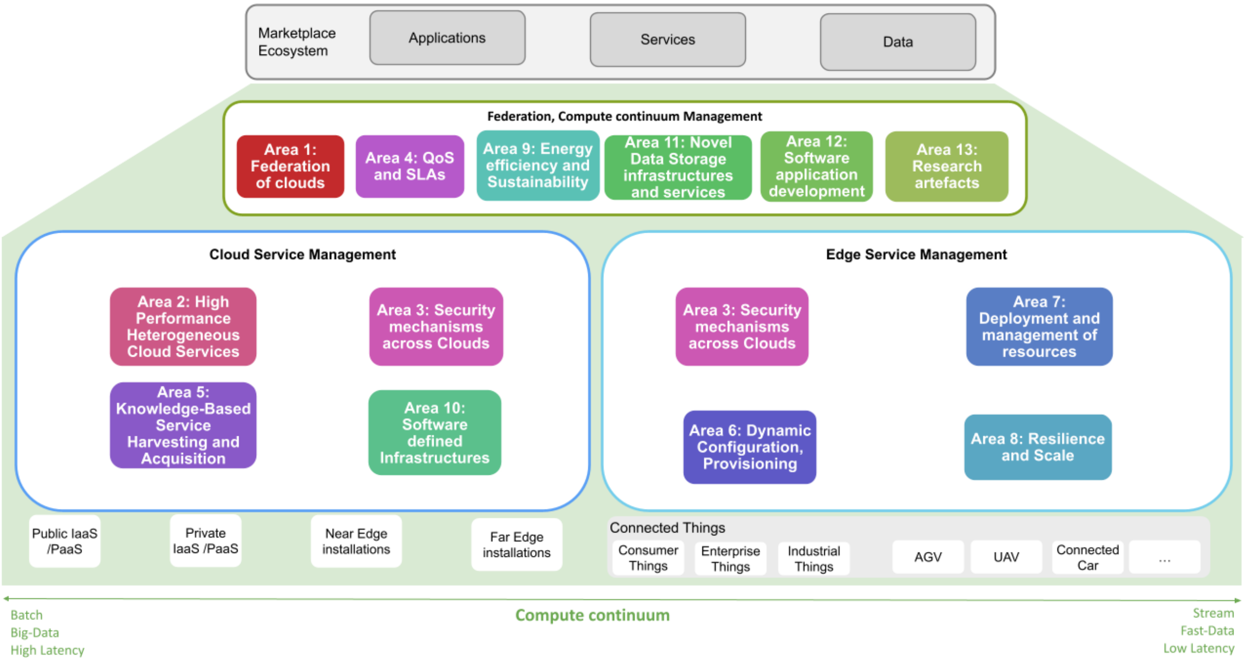
Marketplace Ecosystem
Applications
Services
Data
Federation, Compute continuum Management
Area 1: Federation of clouds
Area 4: QoS and SLAs
Area 9: Energy efficiency and Sustainability
Area 11: Novel Data Storage infrastructures and services
Area 12: Software application development
Area 13: Research artefacts
Cloud Service Management
Area 2: High Performance Heterogeneous Cloud Services
Area 3: Security mechanisms across Clouds
Area 5: Knowledge-Based Service Harvesting and Acquisition
Area 10: Software defined Infrastructures
Edge Service Management
Area 3: Security mechanisms across Clouds
Area 7: Deployment and management of resources
Area 6: Dynamic Configuration, Provisioning
Area 8: Resilience and Scale
Public IaaS /PaaS
Private IaaS /PaaS
Near Edge installations
Far Edge installations
Connected Things
Consumer Things
Enterprise Things
Industrial Things
AGV
UAV
Connected Car
...
Batch
Big-Data
High Latency
Compute continuum
Stream
Fast-Data
Low Latency

Figure 7 Mapping of Future Cluster Research Areas to Architecture levels

- **Research Area 1: Federation of clouds, Facilitate Cloud and Edge interoperability and portability,** taking into account data privacy, security and applicable legislation.
- **Research Area 2: High Performance Heterogeneous Cloud Services,** Management of high-performance computing and high-throughput computing resources
- **Research Area 3: Security mechanisms across Clouds,** Security controls across providers
- **Research Area 4: QoS and SLAs**, Enforcement of quality of service across cloud models

- **Research Area 5: Knowledge-Based Service Harvesting and Acquisition,** Ability to discover and compose services from existing cloud offerings marketplaces
- **Research Area 6: Dynamic Configuration, Provisioning, and Orchestration of Resources,** Orchestration of applications across diverse Edge and Cloud environments
- **Research Area 7: Deployment and management of resources:** in a decentralised, autonomous way Autonomic resource management
- **Research Area 8: Resilience and Scale**, Continuity in service provision for Edge limited operation
- **Research Area 9: Energy efficiency and Sustainability for Edge and Cloud Continuum,** Energy optimisation in Edge and Cloud environments
- **Research Area 10: Software defined Infrastructures and Novel composition models**, Mechanisms for abstraction and virtualisation of resources
- **Research Area 11: Novel Data Storage infrastructures and services,** Strategies for data intensive application execution
- **Research Area 12: Software application development for the computing continuum,** Software engineering, design and programmability of applications in Cloud continuum
- **Research Area 13: Research artefacts for the computing continuum**, Researchers access to tools, data and infrastructures

The identified research areas and challenges are directly related to the three main building blocks we have identified in the Reference architecture presented in Section 4. Figure 7 represents the relation among Research Areas and Reference architecture building blocks and develops the specific relation among these and research challenges. Finally, Table 5 Future Cloud Research Challenges present the relation among the specific future challenges identified in the Future Cluster Research Roadmap and building blocks proposed in the Reference architecture.

Table 5 Future Cloud Research Challenges per layer

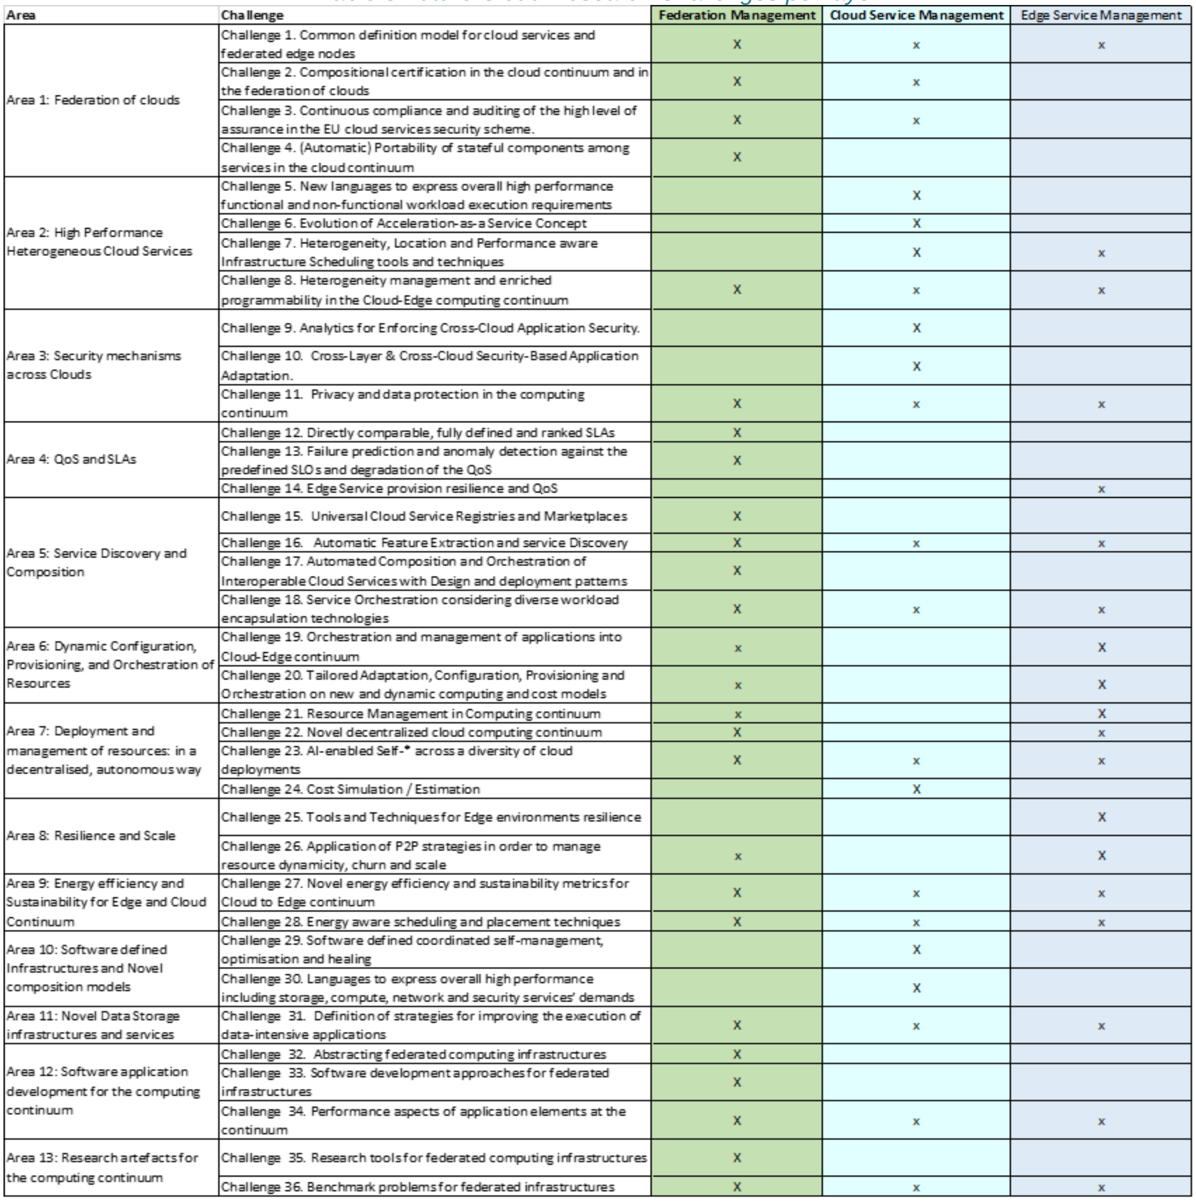

Area	Challenge	Federation Management	Cloud Service Management	Edge Service Management
Area 1: Federation of clouds	Challenge 1. Common definition model for cloud services and federated edge nodes	X	x	x
	Challenge 2. Compositional certification in the cloud continuum and in the federation of clouds	X	x	
	Challenge 3. Continuous compliance and auditing of the high level of assurance in the EU cloud services security scheme.	X	x	
	Challenge 4. (Automatic) Portability of stateful components among services in the cloud continuum	X		
Area 2: High Performance Heterogeneous Cloud Services	Challenge 5. New languages to express overall high performance functional and non-functional workload execution requirements		X	
	Challenge 6. Evolution of Acceleration-as-a Service Concept		X	
	Challenge 7. Heterogeneity, Location and Performance aware Infrastructure Scheduling tools and techniques		X	x
	Challenge 8. Heterogeneity management and enriched programmability in the Cloud-Edge computing continuum	X	x	x
Area 3: Security mechanisms across Clouds	Challenge 9. Analytics for Enforcing Cross-Cloud Application Security.		X	
	Challenge 10. Cross-Layer & Cross-Cloud Security-Based Application Adaptation.		X	
	Challenge 11. Privacy and data protection in the computing continuum	X	x	x
Area 4: QoS and SLAs	Challenge 12. Directly comparable, fully defined and ranked SLAs	X		
	Challenge 13. Failure prediction and anomaly detection against the predefined SLOs and degradation of the QoS	X		
	Challenge 14. Edge Service provision resilience and QoS			x
Area 5: Service Discovery and Composition	Challenge 15. Universal Cloud Service Registries and Marketplaces	X		
	Challenge 16. Automatic Feature Extraction and service Discovery	X	x	x
	Challenge 17. Automated Composition and Orchestration of Interoperable Cloud Services with Design and deployment patterns	X		
	Challenge 18. Service Orchestration considering diverse workload encapsulation technologies	X	x	x
Area 6: Dynamic Configuration, Provisioning, and Orchestration of Resources	Challenge 19. Orchestration and management of applications into Cloud-Edge continuum	x		X
	Challenge 20. Tailored Adaptation, Configuration, Provisioning and Orchestration on new and dynamic computing and cost models	x		X
Area 7: Deployment and management of resources: in a decentralised, autonomous way	Challenge 21. Resource Management in Computing continuum	x		X
	Challenge 22. Novel decentralized cloud computing continuum	X		x
	Challenge 23. AI-enabled Self-* across a diversity of cloud deployments	X	x	x
	Challenge 24. Cost Simulation / Estimation		X	
Area 8: Resilience and Scale	Challenge 25. Tools and Techniques for Edge environments resilience			X
	Challenge 26. Application of P2P strategies in order to manage resource dynamicity, churn and scale	x		X
Area 9: Energy efficiency and Sustainability for Edge and Cloud Continuum	Challenge 27. Novel energy efficiency and sustainability metrics for Cloud to Edge continuum	X	x	x
	Challenge 28. Energy aware scheduling and placement techniques	X	x	x
Area 10: Software defined Infrastructures and Novel composition models	Challenge 29. Software defined coordinated self-management, optimisation and healing		X	
	Challenge 30. Languages to express overall high performance including storage, compute, network and security services' demands		X	
Area 11: Novel Data Storage infrastructures and services	Challenge 31. Definition of strategies for improving the execution of data-intensive applications	X	x	x
Area 12: Software application development for the computing continuum	Challenge 32. Abstracting federated computing infrastructures	X		
	Challenge 33. Software development approaches for federated infrastructures	X		
	Challenge 34. Performance aspects of application elements at the continuum	X	x	x
Area 13: Research artefacts for the computing continuum	Challenge 35. Research tools for federated computing infrastructures	X		
	Challenge 36. Benchmark problems for federated infrastructures	X	x	x

References

[1] The NIST Definition of Cloud Computing, https://csrc.nist.gov/publications/detail/sp/800-145/final

[2] Kenji E. Kushida, Jonathan Murray, and John Zysman. 2015. Cloud computing: From scarcity to abundance. J. IndustryCompet. Trade 15, 1 (2015), 5–19. DOI:https://doi.org/10.1007/s10842-014-0188-y

[3] The NIST Cloud Federation Reference Architecture, https://www.nist.gov/publications/nist-cloud-federation-reference-architecture

[4] Hybrid Cloud Solutions, https://www.vmware.com/cloud-solutions/hybrid-cloud.html#:~:text=VMware%20hybrid%20cloud%20is%20based,cloud%20infrastructure%20in%20the%20world

[5] AWS Outposts, https://aws.amazon.com/outposts/

[6] European Open Science Cloud, https://marketplace.eosc-portal.eu/

[7] GAIA-X: A Federated Data Infrastructure for Europe, https://www.data-infrastructure.eu/GAIAX/Navigation/EN/Home/home.html

[8] Mahadev Satyanarayanan. 2017. The emergence of edge computing. Computer 50, 1 (2017), 30–39. DOI:https://doi.org/10.1109/MC.2017.

[9] Pedro Garcia Lopez. 2015. Edge-centric Computing: Vision and Challenges. ACM SIGCOMM Computer Communication Review September 2015. https://doi.org/10.1145/2831347.2831354

[10] https://www.data-infrastructure.eu/GAIAX/Redaktion/EN/Publications/gaia-x-technical-architecture.pdf?__blob=publicationFile&v=5

[11] NIST Cloud Federated Reference Architecture (CFRA), https://doi.org/10.6028/NIST.SP.500-332

[12] NIST Cloud Federated Reference Architecture (CFRA), https://doi.org/10.6028/NIST.SP.500-332

[13] www.internationaldataspaces.org

[14] https://www.internationaldataspaces.org/wp-content/uploads/2021/01/IDSA-Position-Paper-GAIAX-and-IDS.pdf

[15] Swarm Computing, https://atos.net/wp-content/uploads/2018/12/atos-swarm-computing-white-paper.pdf

[16] https://cloud.vmware.com/vmware-hcx

[17] https://cloud.google.com/anthos

[18] https://www.egi.eu/federation/egi-federated-cloud/

[19] https://egi-federated-cloud-integration.readthedocs.io/en/latest/faq.html#do-i-lose-control-on-who-can-access-my-resources-if-i-join-federated-clou

[20] https://www.openstack.org/

[21] https://opennebula.io/

[22] https://cloud.vmware.com/vmware-hcx

[23] https://cloud.google.com/anthos

[24] https://www.egi.eu/federation/egi-federated-cloud/

[25] HEADS Project, D3.3. Final Framework of resource-constrained devices and networks,http://heads-project.eu/sites/default/files/HEADS%20D3.3%20V1.0.pdf

[26] https://jclouds.apache.org/

[27] https://libcloud.apache.org/

[28] https://github.com/apcera/libretto

[29] http://unikernel.org/

[30] https://katacontainers.io/

[31] https://gvisor.dev/

[32] Zhou, H., Hu, Y., Ouyang, X., Su, J., Koulouzis, S., Laat, C., Zhao, Z.: CloudsStorm: A framework for seamlessly programming and controlling virtual infrastructure functions during the DevOps lifecycle of cloud applications. Softw: Pract Exper. 49, 1421–1447 (2019). https://doi.org/10.1002/spe.2741

[33] Zhou, H., Ouyang, X., Ren, Z., Su, J., de Laat, C., Zhao, Z.: A Blockchain based Witness Model for Trustworthy Cloud Service Level Agreement Enforcement. In: IEEE INFOCOM 2019 - IEEE Conference on Computer Communications. pp. 1567–1575. IEEE, Paris, France (2019). https://doi.org/10.1109/INFOCOM.2019.8737580

[34] Štefanič, P., Cigale, M., Jones, A.C., Knight, L., Taylor, I., Istrate, C., Suciu, G., Ulisses, A., Stankovski, V., Taherizadeh, S., Salado, G.F., Koulouzis, S., Martin, P., Zhao, Z.: SWITCH workbench: A novel approach for the development and deployment of time-critical microservice-based cloud-native applications. Future Generation Computer Systems. 99, 197–212 (2019). https://doi.org/10.1016/j.future.2019.04.008

[35] Switch-2: Zhao, Z., Taal, A., Jones, A., Taylor, I., Stankovski, V., Vega, I.G., Hidalgo, F.J., Suciu, G., Ulisses, A., Ferreira, P., Laat, C. de: A Software Workbench for Interactive, Time Critical and Highly Self-Adaptive Cloud Applications (SWITCH). In: 2015 15th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing. pp. 1181–1184. IEEE, Shenzhen, China (2015). https://doi.org/10.1109/CCGrid.2015.73.