Home-grown, European Cloud and Software Services
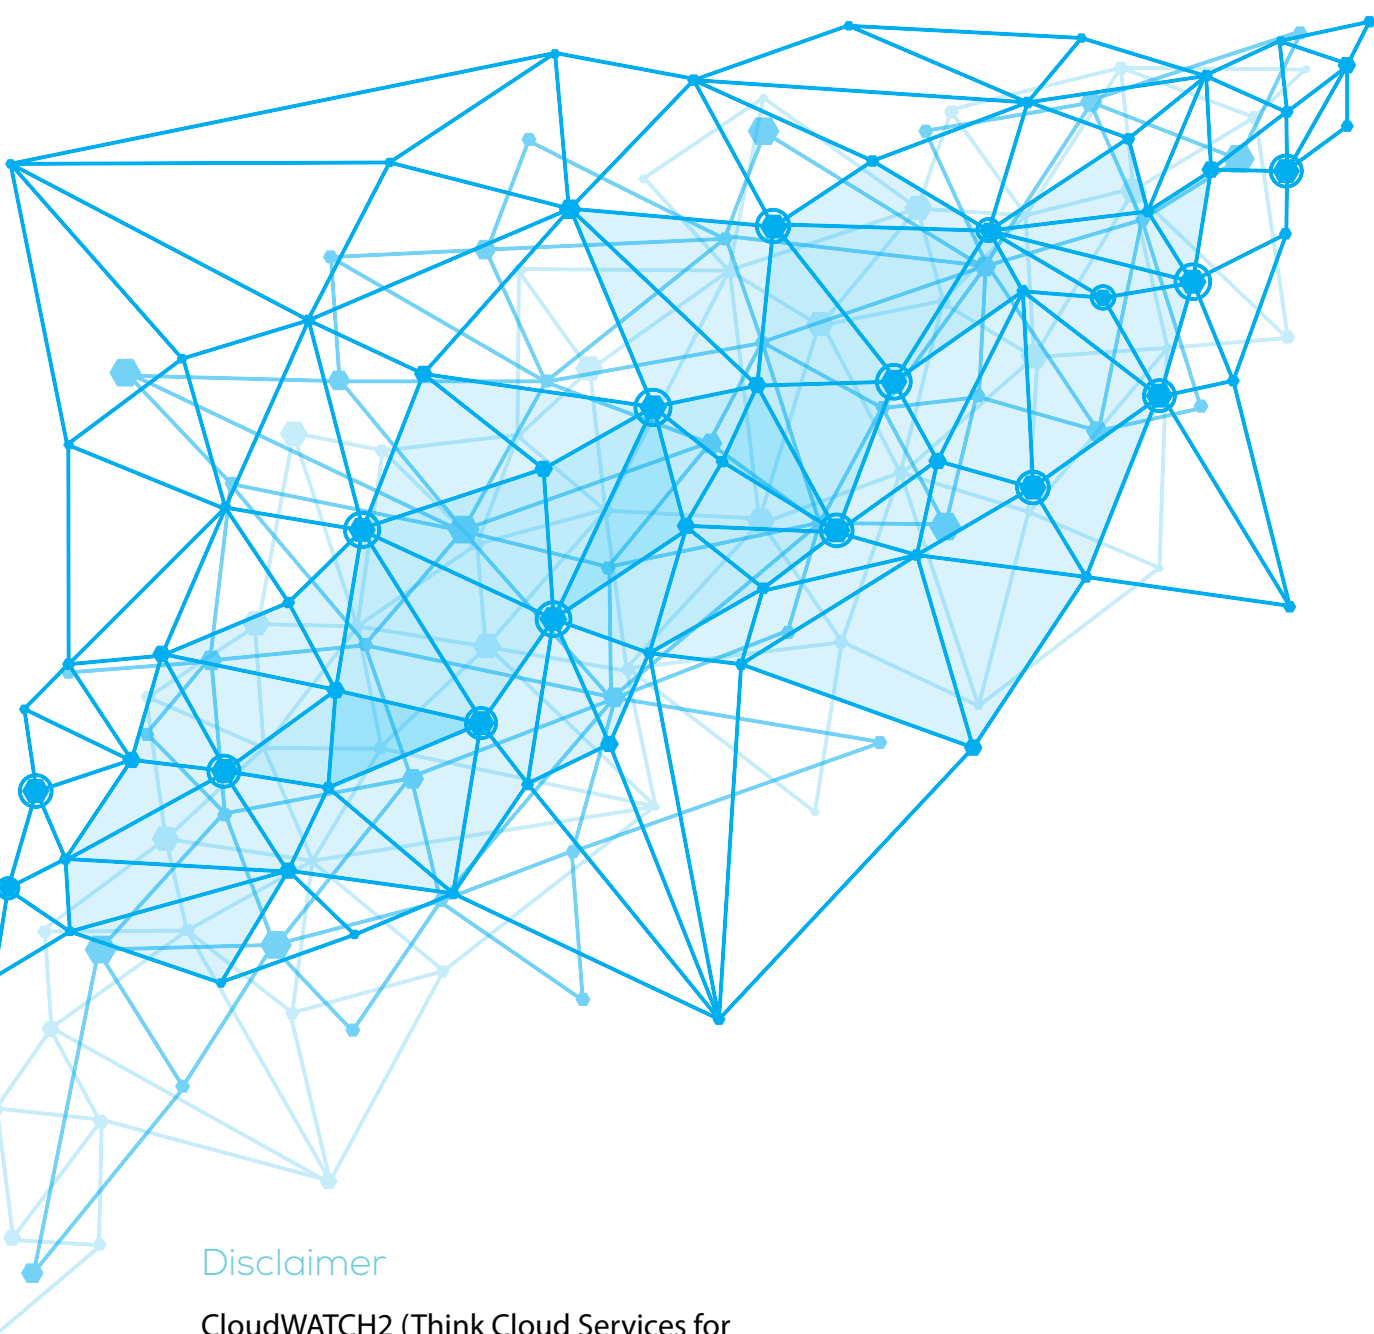
A service catalogue to build trust and security for the Digital Single Market

June 2017

Cloud Watch

Think Cloud Services for Government, Business & Research

Disclaimer

CloudWATCH2 (Think Cloud Services for Government, Business and Research) and the projects included in this publication are funded by the European Commission's Unit on Cloud and Software within DG Connect under Horizon 2020.
The information in this document has been provided by the contributing projects, including CloudWATCH2. Views included in this document cannot be considered to reflect the views of the European Commission.

Summary

A service catalogue to build trust and security for the Digital Single Market

At this the fifth CloudWATCH Concertation meeting, we are delighted to present the second catalogue of offers for trusted and secure services from 54 ongoing projects under the Cloud and Software Unit of DG CONNECT.

The Digital Single Market (DSM) in Europe is an unprecedented opportunity to create one of the biggest digital marketplaces in the world, expected to have an immense impact on Europe's competitiveness and productivity across industrial and service sectors. Cloud computing and software technologies are key enablers across industry sectors, from retail and smart manufacturing to smart agriculture, as well as for Small and Medium Enterprises (SME), large companies and the public sector.

The catalogue demonstrates how European research priorities have adapted to a fast-moving and rapidly-growing market. H2020 Call 1 projects are now providing exploitable results ready for use by SMEs and public administrations in 2017-18. These forward-looking initiatives bring advanced cloud infrastructures and services with improved security, data protection, and quality of service in federated, heterogeneous and multi-layered cloud environments. In addition, new tools and methods for software development are now available to increase productivity in data-intensive systems and distributed applications, as well as tools for managing complex systems.

Looking ahead to 2019, newly-funded projects continue to deliver cloud and software services which are trusted and, with the General Data Protection Regulation fast approaching, secure. Cloud services look towards new paradigms of densely interconnected decentralised cloud infrastructures, with computing and network infrastructures which are tightly interconnected, while also providing robust, trustworthy and high-performance solutions for SMEs and public-sector organisations. From a software perspective, outputs will meet the need in cloud, IoT and big data for programming and modelling methods that facilitate the development of interconnected, flexible, reliable, secure and efficient software.

In addition, the catalogue showcases examples of international collaboration between Europe and Japan, and Europe and Brazil.

Trusted and interoperable services are essential for a thriving Digital Single Market. The exploitation of 'home grown' research and innovation is most timely with the need to remain competitive through technological excellence, sometimes against shrinking budgets, and in view of new opportunities presented by the European Digital Single Market. Projects though face several challenges when defining and implementing their exploitation plans. What they need is a clear understanding of the cloud market and the supply chains within it. Projects not only need to start planning early, but also plot their outputs as the project evolves from a technical viewpoint, making sure they also achieve the right market readiness level.

This catalogue showcases the wealth of expertise and project assets emerging from R&I efforts, and an example of projects stepping-up efforts to actually take results to market. It has been compiled by the CloudWATCH2 project and we'd like to thank all projects for providing their service offers.

More solutions and, software and services are available online:

www.cloudwatchhub.eu/service-offer-catalogue

Concertation meeting in numbers

5th
CloudWATCH Concertation

54
Service Offers

19
New Projects Presented

4
EC Clusters

Market Readiness Assessment of EC funded projects

Embedded within CloudWATCH2's mission to assess and promote EC projects' outputs and results, project partners University of Oxford and Strategic Blue joined efforts with Frank Bennett, iCloud & Cloud Industry Forum to take stock of project assets and potential avenues of commercialisation. The work undertaken in this effort took shape in the form of focused project assessments combined with desktop research and a mix of well-known and emerging business strategy methodologies, for instance business model canvases, value proposition canvases and technology adoption lifecycle.

Over the course of the last year, the team has assessed six EC funded projects (MUSA, CLOUDTEAMS, WAZIUP, CLARUS, EUBra-BIGSEA, COLA) that were all in different phases of their project lifecycle at the time of assessment. While some projects were near to closure with at times having barely more than six months "to go", others were at the height of their productivity and some just starting.

Consequently, the nature of the assessment as well as advice given to these projects were very diverse, though similarities across projects were frequent. For example, projects in later stages of their lifecycle were more interested in specific exploitation advice related to service operation, specific marketing activities, and engagement with different market stakeholders (e.g. visionaries and technology enthusiasts as early adopters in a market segment), etc. Projects at an earlier life stage appreciated suggestions on project execution and leadership and the assessor's input from horizon scanning outside of the project to consider external influences.

However, across all six projects one suggestion stood out and was appreciated by all: To change the public narrative of the project, to shift it from an all too tempting focus on the project's problem space to the intended customer's problem space - while at the same time writing and talking about the same achievements.

This shift of attention was most aptly described by Frank Bennett during a Market Readiness assessment panel discussion, citing "*There are many things 'we could' do with technology, however in a highly competitive market space, the question to ask is whether 'we should' do them.*" This acknowledges the commercial imperative to allocate resources based on an informed understanding of the market potential and is the raison d'etre of the methodology described here.

That statement resonated very well with an observation/statement made by James Mitchell of Strategic Blue when presenting the status of the cloud computing roadmap earlier at the same event: *"Innovation for the sake of innovation is interesting... but pointless. Innovation needs to go to market."*

Both observations share the underlying strong notion of *purpose*: What is the purpose of an innovation and, in the context of EC funded projects, what is the purpose of the project? In order to successfully bring innovation to market it is vital that project partners understand the end-purpose and the Value Proposition Canvas is used to do that.

This is illustrated by a misconception we found common in almost all assessed projects. While all projects were performing notably well with regards to their chosen technological path we found that projects projected the intended use of their innovation as the purpose of the customer's business - and describing this in a language that is different to that used by customers and other entrepreneurs. All assessed projects were familiar with the concept of Value Proposition Canvas, however the language had academic appeal and lacked an understanding of how a customer would apply the innovation in business scenarios in the language of the business' decision team. This had a knock-on effect on how customers' pains and gains were described, more from the viewpoint of the project's output than the value delivered in a business/organisational setting.

In exercising our work we paid due respect to a project's need for confidentiality. Our work with projects is described as amicable and open while recognising there is undeniably a sense of competition and confidentiality was addressed under NDA regime. Any specific findings and suggestions are kept confidential unless the project decides to publicise them.

A summary of the approach used in the CloudWATCH Technical and Market Level Training can be found here.

Advanced Cloud Infrastructures and Services – Exploitable outputs ready in 2017-18

BEACON – Enabling Federated Cloud Networking

Project start: February 2015
Project end: July 2017
www.beacon-project.eu

The increasing interest in hybrid clouds and cloud federations in general raises many challenges in how to deal with the heterogeneity of different cloud platforms. The BEACON solution recommends the **federation of virtual networks** in order to form a federated cloud network that connects all cloud virtual machines that need to interact. The main advantages are that the federated cloud network is isolated from other traffic and is independent of the underlying software and hardware. Another important advantage from the security point of view is that the federated cloud network can be protected with a **specific global network security policy**. The global security policy is enforced with network function virtualisation (NFV) and service function chaining (SFC).

Who is BEACON designed for?

The BEACON solution is designed for user of **private cloud platforms** such as OpenStack and OpenNebula that want to build hybrid clouds and link to **public cloud platforms** such as Amazon WS. More specifically it allows OpenNebula users to build a **federated cloud network for hybrid clouds** connected to Amazon WS for example. With BEACON, OpenStack users can build a federated cloud network for an OpenStack federation.

How will BEACON benefit end-users?

The BEACON solution **automates the management of federated cloud networks**. Federating cloud networks and managing them by hand is a time consuming and error prone activity. BEACON provides a federated cloud network management component that creates and monitors the federated cloud network based on a service manifest that defines the network topology.

How can BEACON improve efficiency and security?

The BEACON solution makes the management of federated cloud networks more efficient. It achieves this by automating federated cloud management. In addition, it makes it easy to define and deploy user defined general network security policies for the federated cloud network.

Cloud Characteristics:

- Advanced security
- Board Network access
- Geographic distribution
- Homogeneity
- Low-cost software
- Massive scale
- Measured service
- On-demand self-service
- Rapid elasticity
- Resilient computing
- Resource pooling
- Service orientation
- Technical
- Virtualisation

CLARUS – A Framework for User Centred Privacy and Security in the Cloud

A framework for user centred privacy and security in the cloud
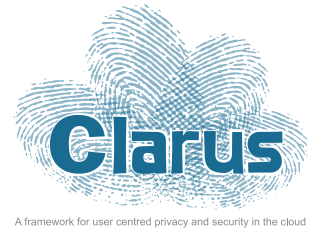
Project start: December 2014

Project end: December 2017

www.clarussecure.eu

The Clarus project provides innovation in **security-enabling techniques, attack-tolerant systems** and in new architectures for secure delivery in the cloud. Its Privacy-preserving mechanisms support data anonymization, data coarsening and splitting, data encryption and searchable encryption. The CLARUS security-by-design approach unleashes new market opportunities by making cloud services more transparent, standardised, auditable and controllable.

Who is the CLARUS designed for?

Information security professionals, and especially **security managers**, can implement privacy-by-design, which ensures users of cloud services are in the driving seat when it comes to controlling their data, such as healthcare IT teams and geospatial communities.

Open source developers, **tech integrators** and **cloud service providers** can support the CLARUS proxy solution and potentially new, secure-by-design services to their customers as the market moves closer towards security-as-a-service. CLARUS can play a key role in not only moving critical applications to the cloud but also in enabling migration from a private cloud service to a pubic one, thus increasing business benefits.

How will CLARUS benefit end-users?

Most security mechanisms are commonly located within the cloud platform. This makes the cloud an impractical solution for those customers whose data is considered sensitive and therefore critical, or for organisations that need to comply with specific regulations on data handling. Cloud customers also need to be assured that no intruder (within or outside an organisation) can hack the cloud and/or impersonate them, and that no denial of service will occur.

With CLARUS, cloud customers no longer need "blind trust" in their cloud service provider(s) when outsourcing their data to the cloud. CLARUS enhances trust in cloud services through its **secure and attack tolerant framework** for the storing and processing of data outsourced to the cloud. This allows end users to **monitor, audit and control** the stored data without impairing the functionality (including the functionality provided by high-level services such as data storage, management, retrieval, transformation, etc.) and cost-saving benefits of cloud services. The attack-tolerant framework is based on a variety of security mechanisms controlled by cloud users without impairing functionalities provided by high-level services, such as data storage, management, retrieval, transformation, etc.) and without reducing the benefits associated with clouds, such as cost savings, ubiquitous access.

How can the CLARUS improve efficiency and security?

CLARUS benefits are clearly demonstrated for **geospatial data** and **digital health** and can be extended to other uses with tight security, privacy and regulatory requirements.

Hospital IT teams can move to the cloud safely knowing that sensitive data is secure and compliant with legal regulations. Public and private organisations operating with geodata can be confident of secure data publication and processing without compromising privacy and control over data.

Cloud Characteristics:

- Advanced security
- Board Network access
- Geographic distribution
- Homogeneity
- Low-cost software
- Massive scale
- Measured service
- On-demand self-service
- Rapid elasticity
- Resilient computing
- Resource pooling
- Technical
- Virtualisation

CLOUDLIGHTNING – Self-Organising, Self-Managing Heterogeneous Cloud

Project start: February 2015
Project end: January 2018
cloudlightning.eu

The CloudLightning system is a new way of **provisioning heterogeneous cloud resources** to deliver cloud services. It provides cloud service providers with an **energy-efficient, scalable architecture** for managing their cloud infrastructures and through the separation of concerns, makes this architecture more accessible to their customers.

Who is CloudLightning designed for?

Tier 1 hyperscale cloud service providers and operators of data centers for cloud services.

How will CloudLightning benefit end-users?

CloudLightning aims to address the problems surrounding **cloud infrastructure management**, **service specification**, and access to and provisioning of cloud resources faced by both consumers and providers of cloud services. Currently, from a cloud service provider perspective, cloud service delivery models typically rely on over-provisioning to guarantee quality of service and accommodate unpredictable peaks in demand.

CloudLightning **resource management** can help **increase output capacity by up to 80%,** without changing the server's energy consumption. Moreover, by using heterogeneous processing resources energy consumption can be further decreased and time to solution can be shortened. This saves time, energy and ultimately costs both for the provider and consumer. On the other hand, when Cloud consumers interface with their chosen cloud service provider, they are tasked with the provisioning, configuring and optimisation of resources to meet their requirements. This requires a commensurate level of IT expertise and knowledge of the provider's infrastructure.

How can CloudLightning improve efficiency and security?

The self-organising, self-managing principles used by CloudLightning provide consumers with a **user-friendly service level interface** to explicitly declare their **requirements for service delivery**. They can do this by specifying their exact workflow, dataset and time deadline. Through the assembly of dynamic resource coalitions, the self-organising and self-managing system automatically and intelligently locates the required resources and chooses the most appropriate configuration to deliver that service, while respecting both the user-level SLA and the business objectives of the cloud provider.

Cloud service providers can enjoy **energy-efficient, scalable management** of their cloud infrastructures and better overall utilisation of service. CloudLightning's general in-service description language is not only easy to implement, but means that the service configuration presented by the self-organising, self-managing system will always be within the constraints of the service provider's resources thereby providing optimal resource allocation and energy consumption. The coalition formation and service deployment processes negate the need to grant users access to manually configure their service description, thus providing an additional layer of security to the provider.

CloudLightning benefits service providers with the freedom to use its resources optimally and frees the consumer of the need and expertise required to configure cloud service delivery.

Cloud Characteristics:
- Board Network access
- Massive scale
- On-demand self-service
- Rapid elasticity
- Resource pooling

CloudSocket – Business and IT-Cloud Alignment using a Smart Socket
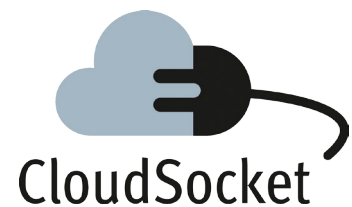
Project start: January 2015
Project end: December 2017
www.cloudsocket.eu

CloudSocket addresses the gap between business and IT worlds, especially for SMEs and provides a decision support for if and how business processes should be moved into the cloud. CloudSocket delivers a **Multi-Cloud enabled business process management (BPM) solution**:

» The CloudSocket BPaaS product enables organisations to **manage their business processes over multiple clouds** by defining, deploying and executing corresponding **self-adaptive workflows** previously designed (by Cloudsocket BPaaS service or other design services/tools).

» The BPaaS consulting service is a set of tools, methods and models that help organisations to decide, if and how to **move their business processes into the cloud**. The tools, methods and models will be used in an interactive manner. In the end, the organisation has enough information to decide whether they are investing in cloud computing or not.

Who is CloudSocket BPaaS designed for?

Potentially any organization with IT – infrastructure, in particular: **Public Administrations** moving towards cloud; **large Companies** that want to continue their cloud strategy also on BP level; **SMEs** who want complete BP solutions rather than several SaaS.

The global Business Process Management (BPM) market is growing. By extrapolating relative BPaaS weight (compared to overall SaaS), as well as BPM market size (as compared to overall software market), we believe that the market for multi-cloud BPaaS product will be around 600 Million USD in 2021.

How will CloudSocket BPaaS benefit end-users?

CloudSocket BPaaS enables organisations to manage their business processes over multiple clouds by defining, deploying and executing corresponding self-adaptive workflows previously designed (by Cloudsocket BPaaS service or other design services/tools).

How can CloudSocket BPaaS improve efficiency and security?

The CloudSocket product intends to close gap between business and the new multi-cloud adoption IT environment by offering an integrated user-friendly solution. Customers could benefit by gaining the ability to efficiently adapt their business needs to current IT proposition i.e. Cloud offerings (to align business processes with Cloud deployed workflows). Thereby, the proposition is to "plug business" into the "Cloud".

BPaaS as a consulting service including special tool support to decide if and how to move business processes into the cloud have to distinctions:

1. USP against other BPM tools, is that those cloud specific analysis mechanisms are yet not implemented in other BPM tools.

2. USP against SaaS, is in its flexibility to quicker establish individual SaaS (BPaaS) offerings and a better integration of services

Cloud Characteristics:

- Advanced security
- Board Network access
- Geographic distribution
- Homogeneity
- Measured service
- On-demand self-service
- Resilient computing
- Resource pooling
- Technical

Vertical Markets:

- Energy
- Engineering & manufacturing
- Finance & insurance
- Smart cities

CloudTeams – Collaborative Software Development Framework based on Trusted, Secure Cloud–based Pool of Users

Project start: March 2015
Project end: February 2017
www.cloudteams.eu

CloudTeams is a cloud-based platform can **transform software development for cloud services** into a much easier, faster and targeted process.

Who is the CloudTEAMS designed for?

The fragmented European market that makes adaptation of innovations difficult and expenses for market research extremely high. In addition, there is a lack of tools and incentives for **users to collaborate with talented software teams** around Europe in order to develop better solutions. CloudTeams targets communities of users who participate in the product life cycle to help software teams develop better solutions for customer problems through team collaboration, customer adaptation and constant testing with trusted users, in a transparent, privacy respectful way.

CloudTeams uses endpoints to existing services and tools that are popular in software development, by mashing them up with common practices. It supports developers with a **collaborative platform** where interaction with customers will feel natural and validates the final outcome. CloudTeams is the intersection of three important fields: **crowdsourcing platforms**, **collaborative software development tools** and **trusted cloud services delivery**. This innovative combination of different tools and practices under a unique concept will be enhanced with a rewarding system to enable end-users to jump into the platform and support solutions they are interested in.

How does CloudTEAMS benefit end-users?

CloudTeams brings together software developers and prospective users, in an unprecedented collaborative platform that **speeds up software development and team collaboration**, by managing and linking information to extract knowledge and predict behaviour of users, to deliver better and valuable solutions.

How can CloudTEAMS improve efficiency and security?

By 2020, CloudTeams aims to be **the prominent tool for trustful software development in Europe**. The CloudTeam Mission is to put customers' needs at the centre of software development, by providing complete software solutions and an active community of prospective users, with respect to their privacy. CloudTeams delivers **a win-win situation for both developers and users** by facilitating relationships of trust where software alleviates pain and be beneficial to the users, while at the same time providing viable business models to the developers.

Cloud Characteristics:

- Advanced security
- Board Network access
- Homogeneity
- Low-cost software
- On-demand self-service
- Resource pooling

Vertical Markets:

- Digital health
- Energy
- Engineering & manufacturing
- Finance & insurance
- Media
- Smart cities

CloudWATCH2 – Think Cloud Services for Government, Business Research

Project start: 30 May 2017
Project end: September 2017
www.cloudwatchhub.eu

The CloudWatchHUB.eu is **the** place to go for **educational guides and practical tools** for the cloud – it's a great place to start for SMEs, schools and small IT teams in public administration. CloudWatchHUB. eu is also the only place with direct access to cutting-edge software services and cloud solutions which are free and available for end- users. Open source is the operative word.

Who is CloudWATCH2 designed for?

Cloud computing is the great leveller offering game-changing capabilities and cost savings for Governments, public administrations, businesses and research. CloudWATCH2 targets these communities to help them make informed decisions when moving to the cloud or taking services to the market, guiding them through the different phases and offering invaluable tips on legal aspects.

How does CloudWATCH2 benefit end-users?

Creating a Digital Single Market comprising a diverse range of digital services underpinned by cloud computing infrastructure, requires a deliberated approach to open source and open standards that enables prospective customers to compare service capabilities and pricing in a transparent way. This is where CloudWatchHUB. eu makes a real difference.

Through our **Cloud Market Roadmap**, we address what a healthy cloud market should look like in order to drive uptake across a very broad set of customers. Research & Innovation outputs are an important part of this and CloudWATCH2 promotes new technologies emerging from European Commission funded projects which provide secure and standardised cloud and software solutions. By assessing **market and technical readiness levels,** CloudWATCH2 provides a valued service to improve the impact of projects. We also **map, test and promote** the adoption of **interoperability and security standards** which are necessary in building trust in the cloud market.

Finally, we educate our stakeholders through **online tools and guides**, and informative **workshops** to promote the adoption of cloud computing.

How can CloudWATCH2 improve efficiency and security?

CloudWATCH2 highlights the importance of the adoption of both standards and benchmarks as being vital if the cloud service market is to become truly trusted. Standards allow a phased approach to cloud adoption using hybrid cloud technologies, and promote vigorous competition across this market. In addition, standards facilitate more accurate **price-performance benchmarking** helping cloud buyers make purchasing decisions.

With more open source standardized solutions tailored for market requirements, R&I project outputs can really contribute to and drive the adoption of cloud computing for government, business and research. With more educated stakeholders, this can lead to a more resilient cloud market with fair, transparent pricing.

CYCLONE – Complete Dynamic Multi–cloud Application Management

Project start: January 2015
Project end: December 2017
www.cyclone-project.eu

Application service providers and developers who deal with complex computing platforms experience a number of deficiencies in the current cloud ecosystem. The CYCLONE solution facilitates the **deployment, management, and use of complex, multi-cloud applications**, as well as enhancing the **end-to-end security** of those applications.

Who is CYCLONE designed for?

The main target are **application developers** or **application service providers** who develop complex computing platforms on cloud infrastructures. The project has identified two flagship applications. An academic cloud platform and associated services for **bioinformatics** research, and a commercial deployment for **smart grids in the energy sector**.

How will CYCLONE benefit end-users?

The current problem while deploying complex applications on federated cloud environment is threefold. Firstly, there is a lack of **automation**. Applications need to be deployed in a distributed way, preparing the application pipeline in advance, checking errors. Developers need to choose where and how to deploy the application and have to secure both the deployment and the running of the application. Secondly, there is a lack of integration with no tool or platform available on the market which covers **application deployment, security and network management** all at once. Finally, a software platform that includes tools for monitoring, scaling, or modifying other aspects of the application is rarely available.

How can CYCLONE improve efficiency and security?

CYCLONE delivers:

» CYCLONE allows users to aggregate cloud resources from both private and public providers to build a cloud platform that is tailored to their application's needs;

» CYCLONE enables dynamic allocation of high-bandwidth channels inside and between data centers;

» CYCLONE allow pluggable monitoring services;

» CYCLONE implements end-to-end cloud security;

Cloud Characteristics:

- Advanced security
- Geographic distribution
- Homogeneity
- Low-cost software
- Massive scale
- Measured service
- On-demand self-service
- Rapid elasticity
- Resilient computing
- Resource pooling
- Service orientation
- Technical
- Virtualisation

Vertical Markets:

- Energy

ENTICE – dEcentralized repositories for traNsparent and efficienT vIrtual maChine opErations

Project start: February 2016
Project end: January 2018
www.entice-project.eu

ENTICE researches and creates a novel VM repository and operational environment for federated Cloud infrastructures aiming to:

» Simplify the creation of lightweight and highly optimised VM images tuned for functional descriptions of applications;

» Automatically decompose and distribute VM images based on multi-objective optimisation and a knowledge base and reasoning infrastructure to meet application runtime requirements;

» Elastic auto-scale applications on Cloud resources based on their fluctuating load with optimised VM interoperability across Cloud infrastructures and without provider lock-in.

Who is ENTICE designed for?

End Users; Cloud Application Providers and/or Software as a Service providers; Application Developers; Cloud Operators and Cloud Providers.

How will ENTICE benefit end-users?

ENTICE use cases from SMEs and industry demonstrate benefits in the areas of energy control and management, earth observation and Cloud orchestration. ENTICE evaluates VM images for size, functionality, and delivery time, where the optimised images should be more than 60% than regular user-created VM images (comprising application users, OS experts and Cloud system administrators) while keeping their original functionality.

Delivery time of the images is expected to be reduced from minutes to around 10 seconds and by 30% for larger images (over 2 minutes delivery time). For VM images in a repository (over 100GB) with over (100 GB) storage requirements will be reduced by 80%. Time for cross data centre deployment reduced by over 20\% compared to manual or semi-automatic techniques with no storage optimisation involved.

The optimisation process should improve the performance of the pilot use cases by at least 30% and decrease costs and storage requirements by at least 25%. The new VM management methods will improve the QoS elasticity of the use cases from their currently inelastic status to elastic.

Finally, the knowledge model will address interoperability and integration issues of the use cases, and will achieve an over 25% of productivity increase in their VM image preparation and deployment time.

How can ENTICE improve efficiency and security?

ENTICE will **reduce the costs** of implementing and using public, private and hybrid cloud infrastructures by **reducing the deployment time of Virtual Machines**, by **reducing the costs of storage in cloud** and by **increasing the performance of the system**, what will lead users to create **real-time scalable** and **elastic applications** at a lower cost. As a middleware, ENTICE allows the user to be abstracted from the infrastructure and implement the applications at a high level through functional descriptors, which makes the system generic to be used independently of the infrastructure used. Finally, ENTICE addresses the needs of **Cloud Application Providers** by providing RESTful application programming interfaces that may be used for dynamic provisioning of highly-optimised VMIs and CIs during the Cloud application runtime.

Cloud Characteristics:

- Geographic distribution
- Massive scale
- Rapid elasticity
- Virtualisation

Vertical Markets:

- Research institutions

ESCUDO-CLOUD – Enforceable Security in the Cloud to Uphold Data Ownership

Project start: May 2017
Project end: December 2017
www.escudo-cloud.eu

ESCUDO-CLOUD invests data owners as first class citizens of the cloud. This will be achieved by providing **enforceable security**, that is, techniques that wrap the data to provide a layer of protection to the storing/processing Cloud Service Provider, setting the trust boundary at the client side, which implies correct and trusted behavior only by the client.

Who is ESCUDO-CLOUD designed for?

ESCUDO-CLOUD is targeted at both **data owners** and **Cloud Service Providers (CSP)**. The availability of technologies enabling users to maintain control over their data can open cloud services to more users or for more applications. It can also free CSPs from the worries of protecting data, allowing them to securely handle the data outside their own control. An additional target are developers of solutions for cloud data management, who can extend their tools with the support of the protection techniques developed in the project.

How will ESCUDO-CLOUD benefit end-users?

ESCUDO-CLOUD provides significant support to **social innovation**. Situations where sensitive information about the user is outside of his/her control will be reduced. The goal is "to **Uphold Data Ownership**," and this aspect is consistent with the evolution of ICT towards a scenario where users are empowered and able to control their own data.

The availability of technologies able to offer marked protection guarantees on outsourced data will thus support a large variety of applications that are becoming increasingly common, thanks to the deployment of a continuously more varied landscape of IT devices.

How can ESCUDO-CLOUD improve efficiency and security?

ESCUDO-CLOUD, will invest users with **control over data** and remove possible concerns that may limit current cloud adoption, enabling users to rely on the cloud with confidence for a variety of applications and data.

ESCUDO-CLOUD will be beneficial to both data owners and CSPs. Data owners will be able to outsource their data while maintaining control over it, with the possibility of regulating access and sharing it with other users in a selective way, and with the insurance that their data will remain protected from the CSPs. Data owners will thus be able to rely on CSPs and use their services for a wider range of applications. This will benefit both companies as well as individual users. Finally, CSPs will significantly benefit from the increased market penetration that robust data ownership would provide, from reduced regulatory risks, audit costs, and general security threats that they would have to face in the absence of such protection.

INPUT – In-Network Programmability for next-generation personal cloUd service supporT

INPUT

Project start: January 2015
Project end: December 2017
www.input-project.eu

The INPUT Project aims to contribute to the evolution of the **Internet "brain"** beyond current limitations due to obsolete IP network paradigms, by moving cloud services much closer to end-users and smart-devices. This evolution will be accomplished by introducing **intelligence and flexibility** ("in-network" programmability) into **network edge devices** and by enabling them to host cloud applications capable of cooperating with and of offloading corresponding applications residing in the users' smart objects and in data centres, in order to realise innovative personal cloud services. This architecture will allow more attention to be placed on main technical challenges like energy efficiency, network programmability, the "softwarization" of the Internet of Things, network and data centre virtualization, personal cloud services, etc.

Who is INPUT designed for?

The Scientific Community (National and European Research & Innovation projects e.g., 5G-PPP) that could utilize the OpenVolcano as beta-users and contribute to its validation in various scenarios/ various cloud services and its roadmap evolution (development of new features) towards commercialisation.

Telecom Operators (mostly European) –upon availability of a commercial release- which could expand their role in the market value chain by offering personal cloud services to individuals (and increase their ARPU); and new advanced hosting capabilities to third-party Service Providers.

How will INPUT benefit end-users?

INPUT's novel infrastructure and paradigm supports Future Internet personal cloud services in a more scalable and sustainable way and with innovative added-value capabilities. INPUT technologies will enable next-generation cloud applications to go beyond classical service models (i.e., IaaS, PaaS, and SaaS), and even to replace physical Smart Devices (SD), usually placed in users' homes (e.g., network-attached storage servers, set-top-boxes, video recorders, home automation control units, etc.) or deployed around for monitoring purposes (e.g., sensors), with their "virtual images," providing them to users "as a Service" (SD as a Service – SDaaS).

INPUT's OpenVolcano platform is a holistic solution enabling the dynamic instantiation into the fog/mobile edge computing facilities of highly demanding personal cloud services by third-parties to end-users with guaranteed QoS/QoE, anywhere, anytime, while ensuring: confidentiality of sensible information; cost efficiency; high manageability and low complexity interfaces; energy efficiency; business viability; easy integration with Service Providers.

How can INPUT improve efficiency and security?

OpenVolcano project is an open-source software platform realizing flexible SDN-driven Personal Networks (rendering end-user devices unnecessary) capable of hosting third-party Personal Cloud Services in fog/mobile edge computing facilities: ease in porting of third-party services benefiting from the existence of well-known state-of-the-art cloud interfaces and APIs (i.e., OpenStack-like); fully 5G-NFV ready; fully autonomous service and infrastructure management; and advanced power management capabilities (compliant to the ETSI GAL v1 and to the upcoming v2 definitions).

Cloud Characteristics:
- Board Network access
- Geographic distribution
- Massive scale
- Virtualisation

IOStack – Software Defined Storage for Big Data

Project start: January 2015
Project end: December 2018
www.iostack.eu

IOStack will enable efficient execution of **virtualized analytics applications over virtualized storage resources** thanks to flexible, automated, and low-cost data management models based on software-defined storage (SDS).

Who is IOStack designed for?

Companies with large datasets of rarely accessed data, system integrators.

How will IOStack benefit end-users?

The IOStack platform **optimises the usage of storage resources** by automatically transferring datasets that are not used as much to a slower but cheaper infrastructure and thus ensuring the optimal usage of all available storage resources.

Our end-users are normally **IaaS providers** or companies operating **big data storage and analytics resources**. They will benefit mainly in simplicity, administration and cost reduction. Instead of operating dedicated clusters, they can virtualize and automate these services thanks to IOStack platform. The advantages of this include high return on investment, cost reduction, and automation.

The IOStack platform will provide a solution in the **OpenStack** platform to the automated management and deployment of virtualized storage and analytic services. Nowadays, the complexity of the tools imply that medium IaaS providers and companies must operate dedicated storage and computing platforms. Or then resort to the major American players in the industry like Amazon or Microsoft. Our solution will **lower the barrier for Big Data storage and analytics** to a variety of small and medium providers.

How can IOStack improve efficiency and security?

The IOStack platform **optimises the usage of storage resources** by automatically transferring datasets that are not used as much to a slower but cheaper infrastructure and thus ensuring the optimal usage of all available storage resources.

Cloud Characteristics:

- Board Network access
- Homogeneity
- Massive scale
- Measured service
- On-demand self-service
- Rapid elasticity
- Resilient computing
- Resource pooling
- Service orientation
- Virtualisation

MIKELANGELO – Improving Responsiveness and Agility of HPC Cloud Infrastructure

MIKELANGELO

Project start: January 2015
Project end: December 2017
www.mikelangelo-project.eu

MIKELANGELO offers flexible, easy to maintain, open source, standardised and user-friendly **Cloud/HPC Cloud technology stack**, focussing on **performance**, **low footprint** and **security**. MIKELANGELO components are easy to deploy and compatible with renowned upstream or enabling open source projects and standards such as KVM, POSIX, OpenStack and Kubernetes.

Who is MIKELANGELO designed for?

MIKELANGELO is addressing basically any **HPC** or **Cloud provider** (public or private), who have identified the need for an IO efficient virtualisation stack, combined with approaches optimised for performance of guest and host execution and communication. HPC virtualisation is impossible without high IO efficiency and very thin guest OS, thus MIKELANGELO is a direct enabler. In the Cloud area, only specific workloads benefit from such efficiency, mainly because IO is of secondary importance in the Cloud. Moving closer to serverless architectures, this will become significantly important for Clouds as well.

How will MIKELANGELO benefit end-users?

MIKELANGELO solves the problem of virtualised IO in HPC and Cloud areas. More importantly, MIKELANGELO is offering a **widely compatible** and **user-friendly set of approaches and tools**, that helps with the **IO optimisation**, independently of the target infrastructure.

With MIKELANGELO, we are employing a **mix and match approach**, where you can take specific components and integrate them into your specific infrastructure. For example, we offer **unikernel** technology (POSIX compliant OSv), with additional set-up tools; We offer **IO optimisation** technologies on the KVM (host) level - ZeCoRX and IOcm components and also between guest and host, with virtual RDMA. These components can be used on their own or all together. The same approach is with their integration - we can use them as-is, or integrate them into e.g. OpenStack or use them in GCE or AWS. Finally, we offer strengthened host security (SCAM component). Using MIKELANGELO components enables **new virtualised architecture approaches** and **increased performance and scalability** of the infrastructure as a whole. MIKELANGELO is also offering a range of supporting approaches and technologies, making the integration and use very straightforward and efficient.

How can MIKELANGELO improve efficiency and security?

MIKELANGELO consists of a set of technologies and approaches, all aiming to improve IO performance or to help with integration of MIKELANGELO components into larger systems. With MIKELANGELO we can use **unikernel** instead of traditional guest operating system in virtualisation. This improves the efficiency, as unikernel is typically very small (50 megabytes), can be loaded in under a second and executes the workload as efficiently as standard guest (eg. Linux). On the host level, using approaches as **IOcm and ZeCoRX** (KVM modules), we can control the **IO performance** of the overall virtualised infrastructure on that specific host. We can even improve the **guest-host communication** or **guest-guest communication** under the same host, using the virtualised RDMA approach (if hardware support is available). It is important to note that OSv unikernel offers POSIX compatibility, however when such compatibility is not required, we can use novel programming approach called Seastar. Using such approach, we can obtain very high-performance improvements, however we need to rewrite the application from scratch. Other approaches in MIKELANGELO are **Scotty**, for scientific experiments deployment and control; **snap** for highly efficient logging and telemetry; **SCAM** for strengthened host security; and **vTorque** for deployment of virtual machines under Torque system, etc.

Cloud Characteristics:

- Advanced security
- Low-cost software
- Massive scale
- Measured service
- Rapid elasticity
- Service orientation
- Virtualisation

Vertical Markets:

- Engineering & manufacturing

MUSA – Multi-cloud Secure Applications

Project start: January 2014
Project end: December 2017
www.tut.fi/musa-project

The The MUSA project provides an integrated tool framework for **DevOps** and **Agile engineering** of **multi-cloud based applications**, addressing **security** in all its phases: design, deployment and operation. The framework supports risk analysis and selection of secure cloud services, and is able to automatically deploy and monitor distributed components and create an application Service Level Agreement.

Who is MUSA designed for?

The main targeted users are **DevOps teams** covering four main roles that need tools that can better integrate a seamless assurance of security in the applications:

» **Application developers** (including architects) require tools to design multi-cloud applications according to functional and, data security features, and security at runtime mechanisms;

» **System operators** require tools to automatically set and deploy best combinations based on functional and security needs;

» **Service administrators** need to monitor the operation of the application (fulfilment of SLA), in order to promptly react to security incidents and inform users;

» **Business managers** responsible for the business aspects of cloud service provision.

How will MUSA benefit end-users?

MUSA supports the control of security in distributed applications over heterogeneous cloud resources. **It reduces time-to-market** and shortens the gap between the Development and Operations for a **timely reaction to security incidents at runtime**. Users access a unique Kanban-style Dashboard that is able to encompass a number of tools that can also be used separately.

The security framework integrates security assurance in both engineering and operation and eases:

» Multi-disciplinary Risk analysis identifying security controls in the application components;

» Selection of cloud services, taking into account security controls on offer;

» Automated creation of SLA requirements of the application and distributed components;

» Automated deployment of distributed components in heterogeneous cloud services;

» Automated monitoring and enforcement of security behaviour granted in the SLA through the use of agents within the application components.

How can MUSA improve efficiency and security?

The data security incidents in multi-cloud applications will be **reduced** through the assurance of a secure behaviour of individual cloud-based components and the overall application, even if the data is processed and/or stored by untrustworthy or opaque cloud providers. The **cloud consumers' trust** in clouds will be **enhanced** by the provision of tools for expressing their security needs and keeping them informed on the security and performance faults of the multiple cloud services in use. **Application developers** will be able to model the **multi-cloud application**, based on the functional and security features on offer in the SLA, as well as to embed application component mechanisms to enforce security at runtime. **System operators** will be able to automatically discover and select the best cloud service combinations by balancing performance and security. **Service administrators** can assure the secure behaviour of multi-cloud applications and minimize the security risks while keeping the users informed. **Business managers** will be able to make better-informed decisions when selecting cloud services.

Cloud Characteristics:
- Advanced security
- Geographic distribution
- Resilient computing

Vertical Markets:
- Digital health
- Finance & insurance
- Smart cities

PaaSword – A Holistic Data Privacy And Security By Design Platform-As-A Service Framework

PaaSword

Project start: January 2015
Project end: December 2017
www.paasword.eu

The PaaSword project introduces a novel **data privacy** and **security by design** framework with the objective of protecting sensitive data stored in the cloud. PaaSword enables security annotations, transparently through an IDE, transformed into context-aware security policies that enforce access control, cryptographic protection and physical distribution to secure the privacy of sensitive data.

Who is PaaSword designed for?

PaaSword extends the Cloud Security Alliance's cloud security principles by capitalising on recent innovations in virtual database middleware technologies, which introduce a **scalable and secure cloud database abstraction layer** with sophisticated data distribution and encryption methods. PaaSword provides **encrypted and distributed storage**, and **context-aware access control**. constituting a valuable asset for any Platform-as-a-Service provider.

How will PaaSword benefit end-users?

The adoption of PaaSword brings about new access control mechanisms that incorporate dynamically changing contextual information into access control policies and context-dependent access rights, which along with the encrypted and distributed storage support, consolidate the perfect fit for the dynamic cloud computing environment. It brings all adopters one step closer to compliance with very demanding security regulations, such as the EC's General Data Protection Regulation. In brief it offers:

» A searchable encryption scheme for secure queries support.

» Policy-based access control & context-aware security models.

» Governance capabilities for ensuring the validity of access control policies.

» A dedicated IDE plug-in for injecting code-level annotations that associates these policies with methods that provide access to sensitive data.

» A novel policy enforcement middleware that extends the well-known attribute-based access control paradigm with semantically-rich context information.

» Unique distributed storage across IaaS providers for disentangling data objects that might reveal sensitive

information to internal or external adversaries.

» The PaaSword holistic framework that integrates all of these novel offerings.

How can PaaSword improve efficiency and security?

Current cloud applications and storage volumes often leave information at risk to theft, unauthorized exposure or malicious manipulation. The most critical part of a modern cloud application and services is the data persistency layer and the database itself. PaaSword introduces a **holistic data privacy** and **security by design framework**, based on distributed and encrypted data persistence and sophisticated context-aware access control mechanisms in cloud-based services and applications.

Unlike any other solution, PaaSword supports both developers of cloud applications with code annotation techniques and DevOps with the necessary modelling and management tools for achieving an appropriate level of protection for their cloud application's data, even in cases where sensitive information resides on untrusted IaaS providers. Thus, PaaSword enables enterprises to unlock the valuable business, economic and operational benefits of migrating to the cloud, as it generates the confidence of individuals and corporate customers in cloud-enabled services and applications. These valuable business benefits cannot be unlocked without addressing the new data security challenges posed by cloud computing.

The long-term expectation of the impact of the project is to assist in the accelerated adoption of cloud computing technologies, and to see a paradigm shift of European industry towards security and privacy.

Cloud Characteristics:

❯ Advanced security
❯ On-demand self-service
❯ Resource pooling
❯ Service orientation

Vertical Markets:

❯ International agencies
❯ Research institutions

RAPID – Heterogeneous Secure Multi–level Remote Acceleration Service for Low–Power Integrated Systems and Devices

Project start: January 2015
Project end: December 2017
rapid-project.eu

RAPID is the first public acceleration cloud service. The secure unified model lets almost any device or infrastructure operate as an accelerated entity and/or as an accelerator serving other less powerful devices in a secure way. This ranges from smartphone, notebook, laptop and desktop to private and public cloud can RAPID also offers a registration mechanism, which permits the accelerated entities to automatically find and connect to nearby accelerators with the required resources. The RAPID platform will become available and commercially exploitable being offered as a service on a pay-per-use basis or under a subscription fee, depending on the user profile (accelerator or accelerated entity).

RAPID targets a novel heterogeneous CPU-GPU multi-level cloud acceleration focusing on applications running on embedded systems found on low-power devices such as smartphones, notebooks, tablets, wearable devices (smart watches, glasses), robots, cars, etc. Many such low-power devices can't always cope with the increased demand for processing power, memory and storage required by several applications in entertainment, vision, security, robotics, and aerospace such as gaming, antivirus, augmented reality, face and speech recognition, movement detection, biometrics, and CCTV. These applications require tremendous performance and cannot run on most existing battery-operated devices. As a result, most such applications are only executed on high-end servers.

Who is RAPID designed for?

Many low-power devices such as smartphones, tablets, notebooks as well as several other embedded systems can't always cope with the increased demand for processing power, memory and storage required by modern applications in gaming, vision, security, robotics, aerospace, etc. As a result, most such applications are only executed on high-end servers. RAPID tackles this challenge by taking advantage of high-performance accelerators and high-bandwidth networks. RAPID will provide a direct impact to different stakeholders, mainly end user, application developers and cloud providers (IaaS, PaaS and SaaS).

How will RAPID benefit end-users?

RAPID allows end users to register in order to allow its accelerated entity (low power device) to automatically **find and connect the nearby accelerators**. It provides an expressive programming model which abstracts to developers the complexity of managing the acceleration mechanism. For **cloud providers**, RAPID provides means to **improve performance** while taking care of **QoE** aspects, a distributed heterogeneous infrastructure that can change the future of mobile applications and will create new innovation opportunities to service providers by introducing Acceleration as a Service.

How can RAPID improve efficiency and security?

» RAPID will accelerate the use of application on almost any end user device.

» RAPID supports developers creating "accelerated" applications for customers.

» RAPID will improve operational expenditure of IaaS providers.

» RAPID supports PaaS providers to deliver accelerated services to developers.

» RAPID enriches SaaS applications executable in low-power devices.

Cloud Characteristics:

- Advanced security
- Board Network access
- Low-cost software
- Measured service
- On-demand self-service
- Rapid elasticity
- Resource pooling
- Service orientation
- Technical

Vertical Markets:

- Digital health
- Media
- Smart cities

SERECA – Secure Enclaves for REactive Cloud Applications

Project start: March 2015
Project end: March 2018
www.serecaproject.eu

The SERECA project provides an infrastructure that **protects the integrity and confidentiality of applications** running in **untrusted cloud environments**, even against malicious insiders with super-privileges. This relies on a set of facilities/tools useful to transparently leverage the new ISA extension of Intel's CPU, namely Software Guard eXtension (SGX).

Who is SERECA designed for?

SERECA targets organisations/companies that are willing to migrate to the cloud, but are still unhappy with the security guarantees of current cloud offerings e.g. the Critical Infrastructure (CI) sector. SERECA demonstrates that the secure mechanisms provided by the platform fully satisfy the stringent security requirements of a CI. In addition, SERECA tools are readily available to a wide class of distributed applications in a cloud environment and requiring superior security.

SERECA's **RiskBuster** pilot application will monitor assets of a civil water supply network. Many sensors distributed on a dam will provide sensitive data that need to be handled carefully by the application. The **Illuminate** pilot application is an Application Performance Management (APM) SaaS solution. Security is increased allowing Illuminate to store, retrieve and process sensitive data used in the service. It is also securely hosted on multiple cloud providers.

How will SERECA benefit end-users?

SERECA advances the state-of-the-art of cloud technology for **confidentiality** and **integrity**. By protecting applications running on the cloud from malicious insiders, e.g.: provider employee or contractor SERECA enhances security levels and reduces how much the user need to trust their cloud provider. SERECA also provides protection against **privileged software** (e.g. the hypervisor), which can also be exploited by attackers to get access to virtual machine data. The SERECA infrastructure protects sensitive data building on hardware extension of new Intel's CPUs, which allow critical operations to be executed in a protected area of memory, called secure enclave, that is inaccessible even to users with higher privileges.

How can SERECA improve efficiency and security?

SERECA enables secure processing of sensitive data in a protected area of memory, which is inaccessible even to users with super-privileges. **SERECA secure execution environment** (namely: Secure Container) provides applications with tool for handling critical data. Data can be in encrypted form at all times, except when it is under control of the processor for actual usage. The SERECA platform keeps all **in-memory data encrypted** and only the application itself has access to the memory. Keys are kept confidential: only the application can access its keys, neither the root user on the VM nor the root user on the physical machine has access to these keys. Beside protecting data processing, the SERECA secure communication mechanism (called Secure Bus) also performs protection when data travels between the cloud environment and the Internet. Moreover, a **Secure Coordination Service** is provided in SERECA to run the secured applications on a distributed platform, for better reliability and performance. Finally, SERECA's facilities also include a **Partitioning Tool**, which enables easy porting of applications including legacy ones and ultimately enable seamless migration to the SERECA platform.

Cloud Characteristics:
- Advanced security
- Board Network access
- Geographic distribution
- Homogeneity
- Massive scale
- Resilient computing
- Resource pooling
- Service orientation
- Virtualisation

Vertical Markets:
- Energy

SSICLOPS – Scalable and Secure Infrastructures for Cloud Operations

Project start: February 2015
Project end: January 2018
ssiclops.eu

The SSICLOPS project will focus on techniques for the **management of federated private cloud infrastructures**, in particular **cloud networking techniques**. Key deliverables from the project will include a meta data description language for workloads, resources and policies, workload-specific adaptations to TCP/IP stacks, and data center performance analysis tools.

Who is SSICLOPS designed for?

Cloud operators, and organisations using cloud services, who care about **performance and privacy** considerations of their cloud services.

Multiple application areas can benefit from SSICLOPS results. The project target scenarios include **in-memory database management systems**, **large-volume high-energy physics computation**, **network function virtualisation** for telecommunication operators and **content distribution networks**.

How will SSICLOPS benefit the end-user?

SSICLOPS develops a **federated cloud infrastructure** with following properties:

» Policy language to describe constraints in federated cloud communication.

» Improved communication efficiency, for example using multipath TCP between federated cloud instances.

» Enhanced protocol stack designs for minimum latency.

» Advanced analysis tools for diagnosing communication performance.

How can SSICLOPS improve efficiency and security?

End users can seamlessly use high-performance federated cloud services with guarantees on specific policies. The project is also involved in **protocol development** and **standardisation**, to make sustaining impact on future protocol development. The protocol stack enhancements are made available as open source, for the benefit of cloud networking community at large.

Cloud Characteristics:

❯ Board Network access
❯ Geographic distribution
❯ Measured service
❯ Resource pooling

Vertical Markets:

❯ Engineering & manufacturing

SUNFISH – SecUre iNFormatIon Sharing in Federated Heterogeneous private clouds

Project start: January 2015
Project end: December 2017
www.sunfishproject.eu

SUNFISH addresses the lack of infrastructure and technology allowing Public Sector players to federate their private clouds and at the same time respect legislative and security barriers to using commercial technological solutions. In particular, SUNFISH's "**Federation-as-a-Service**" approach covers: Federation Administration; Runtime and Offline Monitoring; Data Security Assurance; Data Transformation Services.

Who is SUNFISH designed for?

The SUNFISH platform, a solution tailored specifically for **public sector bodies** and potentially private sector players. Currently there are three real-world use cases implemented: the Italian Ministry of Economy and Finance (MEF), the Maltese Ministry for Finance (MFIN), the South East Regional Organised Crime Unit (SEROCU).

How will SUNFISH benefit end-users?

Nowadays, the privacy and control of information propagation are becoming more and more relevant issues in the public sector. Through these new technologies, the SUNFISH consortium aims at improving security in federated "**national**", as well as "**cross-border**", clouds. The impact of the wider adoption of cloud federation technologies will enable a greater infrastructure usage efficiency, encouraging a better resource utilization of the cloud infrastructures of the Public Administration bodies, thanks to a more effective workload management between shared private clouds.

During the project, a secure approach for federating private clouds is being developed, in order to guarantee a high level of safety for EU citizens and businesses who benefit from Public Services. In addition, a continuous monitoring of inter-cloud communications will be provided, as well as the ability to share services between different private clouds, in a fast, flexible and secure way with reduced management costs of IT infrastructure for Public Sector Entities.

How can SUNFISH improve efficiency and security?

Designed for the public sector, "Federation as a Service (FaaS)" is SUNFISH's unique selling point. This assures dynamic and secure management of cloud federations including **democratic governance**; **dynamic federation**; and **innovative privacy-preserving services**. This reduces the **deployment** and **operating expenditure** of cloud environments in particular, service failure, data privacy and cyber risks. Furthermore, SUNFISH fosters a principled first-time advanced exploitation of **blockchain** technology for realising the federation's registry, managing access control policies, monitoring data communication and storing sensitive data items.

Finally, SUNFISH's deployment leads to lower energy use and consequently lower CO_2 emissions, due to higher efficiency of services provided by the SUNFISH platform.

Cloud Characteristics:

- Advanced security
- Board Network access
- Geographic distribution
- Homogeneity
- Low-cost software
- Massive scale
- Measured service
- On-demand self-service
- Rapid elasticity
- Resilient computing
- Resource pooling
- Service orientation
- Virtualisation

Home-grown, European Cloud and Software Services

SUPERCLOUD – User–centric management of security and dependability in clouds of clouds

Project start: January 2015
Project end: January 2018
www.supercloud-project.eu

Distributed clouds are about systems that should work across multiple administrative domains, some of which may not be trusted. Making user-centric security and dependability for clouds-of-clouds a reality means overcoming 4 challenges. First, **security** should be **self-serviced**, so that users may define fine-grained security settings to control the protection level of their cloud resources. Second, security should be **self-managed** to tackle complexity through automation in infrastructure layers, and across provider domains. Third, security should be **end-to-end**, to define and enforce security SLAs such as for isolation across interoperability barriers for computing and networking. Finally, **resilience** should be guaranteed in a multi-provider setting, for instance in terms of data availability, and in network data and control planes.

SUPERCLOUD proposes new security and dependability infrastructure management paradigms that are:

» User-centric, for self-service clouds-of-clouds where customers define their own protection requirements and avoid lock-ins.

» Self-managed, for self-protecting clouds-of-clouds that reduce administration complexity through security automation.

Who is SUPERCLOUD designed for?

The target is clearly the customer! SUPERCLOUD offers business opportunities in many dimensions, such as the creation of value-added services that bring together resources from several, possibly untrusted, cloud providers to give users better service and more security and dependability guarantees. This means that the customer can choose which security and availability services to deploy in his own self-service cloud. This is illustrated by deployment of the infrastructure for healthcare use-cases. Other areas include include cloud brokerage, network function virtualization, blockchain, or smart home security.

How will SUPERCLOUD benefit end-users?

SUPERCLOUD will build a **self-management infrastructure** for security and dependability of heterogeneous resources across federated clouds. Customers will be provided with self-service

environments enabling adaptive, **customizable security** for their cloud applications and services. SUPERCLOUD will provide innovative **cryptographic** methods and tools for protecting data across distributed clouds through on-demand data security services, such as access control, blind computation, privacy-preserving indexing, and data availability. SUPERCLOUD will enable **resilient network-as-a-service**, leveraging software-defined networking paradigms. It will also provide strong guarantees for end-to-end security and integrated trust management across multiple infrastructure layers and cloud domains.

How can SUPERCLOUD improve efficiency and security?

SUPERCLOUD delivers:

» **Self-Managed Security**: Development of an autonomic security management framework that operates seamlessly over compute, storage and network layers, and across provider domains to ensure compliance with security policies.

» **End-to-End Security:** Proposition of trust models and security mechanisms that enable composition of services and trust statements across different administrative provider domains.

» **Resilience:** Implementation of a resource management framework that composes provider-agnostic resources in a robust manner using primitives from diverse cloud providers.

Cloud Characteristics:

Advanced security
Geographic distribution
Homogeneity
On-demand self-service
Resilient computing
Service orientation

Vertical Markets:

Digital health
Finance & insurance

Tools and Methods for Software Development - Exploitable outputs ready in 2017-18

ALIGNED – Quality-centric Software and Data Engineering

Project start: February 2015
Project end: January 2018
aligned-project.eu

ALIGNED will manage the volume, variety and velocity of data available through the web, with particular focus on linked data, and integrating it into business processes, with particular focus on software engineering.

Who is ALIGNED designed for?

Small & medium enterprises. Large companies. Technology providers, Research institutions.

How will ALIGNED benefit end-users?

Spinning off technology company. Spinning off academic foundation. Injecting innovations into production tools of industrial partners.

The ALIGNED software and data engineering methodology will provide a basis for **unified management** views of co-evolving software and data systems. **Tool interoperability** will be increased by publishing and consuming enterprise Linked Data that describes the lifecycle of the software and data engineering process. In addition, maintaining and integrating software and web data sources will be easier and the resultant quality and agility of these systems will be increased.

How can ALIGNED improve efficiency and security?

ALIGNED improves the ability to **create** and **analyse** high-quality datasets and **integrate** them with business processes more efficiently.

It also enhances engineering processes, methods and tools for data-intensive systems. This includes the development of the **PoolParty** semantic middleware, linked data quality management tools such as RDFUnit and Dacura, the open suite of ALIGNED vocabularies for describing data intensive system engineering and enabling tool-chain integration and model-driven engineering tools such as Booster and the Model catalogue. The Wolters Kluwer JURION Legal Information system will also be enhanced, DBpedia improved in terms of quality and release efficiency and the Seshat: Global History Databank will release expert curated social science data onto the web.

Cloud Characteristics:
- Massive scale
- Service orientation

Vertical Markets:
- Digital health
- Media

ARCADIA – A novel reconfigurable by design highly distributed applications development paradigm over programmable infrastructure

Project start: January, 2015
Project end: December, 2017
www.arcadia-framework.eu

The ARCADIA project aims to design and validate a Novel Reconfigurable-By-Design Highly Distributed Applications Development Paradigm over Programmable Infrastructure. The proposed framework will rely on the development of an extensible Context Model which will be used by developers directly at the source-code level, assisted and validated by IDE-plugins in order to re-assure that the generated executable files contain meaningful semantics. According to ARCADIA's vision, the generated executables should be on-boarded by a Smart Controller which will undertake the tasks of translating annotations to optimal infrastructural configuration. Such a controller will enforce an optimal configuration to the registered programmable resources and will pro-actively adjust the configuration plan based on the Infrastructural State and the Application State.

Who is ARCADIA designed for?

Small & medium enterprises, Large companies, Technology providers, Research institutions.

How will ARCADIA benefit the end-user?

ARCADIA's reference implementation and developed toolkits are going to facilitate application developers to design and develop infrastructural-agnostic applications and lead to the evolvement of novel and innovative paradigms for the deployment of advanced applications, boosting in this way the competitiveness of the software development industry. Based on the use of a set of annotations in the source code level, application deployment and management will be supported through the design and implementation of Smart Controllers.

The Smart Controllers will orchestrate the dynamic allocation and management of resources (e.g. physical devices, virtual images, linux containers), as well as the configuration of the execution environment in real time. The approach followed ensures the optimal use of the available resources based on the existing policies, as well as the optimal runtime configuration in application and infrastructural level. This is based on orchestration provided by distributed collaboration among the Smart Controllers. Thus, optimisation of the application's execution according to multiple objectives (e.g. energy efficiency, QoS, security), as set by the end users, will be supported.

How can ARCADIA improve efficiency and security?

ARCADIA mainly proposes a software development paradigm along with a deployment framework that enables the design, development and deployment of distributed applications over programmable infrastructure. A set of challenges are going to be tackled including (i) the design of a novel software development paradigm that will facilitate software developers in developing applications that can be represented in the form of a service graph that can be deployable and orchestratable over programmable infrastructure, (ii) the design and implementation of an orchestrator (called Smart Controller in ARCADIA) able to undertake the developed service graph and proceed to optimal deployment and orchestration of the corresponding service/application and (iii) the design and implementation of a policy management framework that supports the definition of high and low level policies on behalf of a Services Provider that can be associated with a service graph or a set of service graphs along with the definition of a set of metrics for policies management purposes.

Home-grown, European Cloud and Software Services

CHOReVOLUTION – The Choreography Platform for Future Internet Applications

Project start: January 2015
Project end: December 2017
www.chorevolution.eu

The CHOReVOLUTION platform is solving the problems of **digital service ecosystems and IoT environments** where different actors expose their services. It offers benefits in **smart cities, smart mobility and tourism, B2B applications, smart manufacturing and Business processes automation**. Developing multi-service architectures will become typical in the future IoT-flooded world. Service architectures will become more and more fragmented and things will start to have a substantial computing power and capabilities, making the orchestration paradigm very hard to use. Due to the fact that the CHOReVOLUTION platform enables web services to interact with each other, developers are freed from having to hardcode the web services to first interact with the operator to exchange information with each other. This way the web services are able to dynamically interact with each other to extract data that can be presented to the end-user.

Who is CHOReVOLUTION designed for?

Designers of complex systems can focus on the business aspects of their solution, while delegating the deployment and runtime technicalities. The CHOReVOLUTION platform eases the development of service choreographies by allowing to reduce the time to market and enhance correctness of the built choreographies.

Cloud Service Providers can host the software allowing the interactions between existing or new services with the objective to build a new system. The CHOReVOLUTION platform provides tools to manage these interactions and to generate the software allowing them. The deployed software allows dynamic evolutions of the system.

How will CHOReVOLUTION benefit end-users?

The CHOReVOLUTION platform automates the implementation of a service choreography involving heterogeneous services and things. It provides the tools to implement a secure choreography from the design phase to the cloud deployment. Moreover, the platform also offers web tools to monitor and manage deployed choreography/services in a centralized way.

How can CHOReVOLUTION improve efficiency and security?

The CHOReVOLUTION platform helps innovate with IoT-based services and compound web services more rapidly, more securely and cost-efficiently. It simplifies operation and management tasks, while enabling secure service composition. There is no need to maintain an operator that leads the orchestration. If one module goes offline the rest of the system can replace it and keep operating. It enables an easier way to scale the application developed on the platform.

The innovative aspect of the CHOReVOLUTION platform concerns a great deal with the platforms ability to manage heterogeneous services and things, and to provide automated support for the development, enactment and execution of the specified choreographies.

Cloud Characteristics:
- Advanced security
- Board Network access
- Massive scale
- Measured service
- On-demand self-service
- Rapid elasticity
- Resource pooling
- Service orientation

Vertical Markets:
- Digital health
- Energy
- Engineering & manufacturing
- Finance & insurance
- Smart cities

DICE – Developing Data–Intensive Cloud Applications with Iterative Quality Enhancements

Project start: February 2015
Project end: January 2018
www.dice-h2020.eu

DICE will deliver innovative development methods and tools to strengthen the competitiveness of small and medium Independent Software Vendors in the market of business-critical data-intensive applications (DIA). Leveraging the DevOps paradigm and the innovative Big Data technologies of nowadays, DICE will offer a complete open source solution for the Quality-Driven Development of cloud-based DIA.

Who is DICE designed for?

» **Large IT-based organisations** with multiple market segments and multiple DIA needs.

» **Big Data Infrastructure Providers** (Cloud services providers, Archiving Solution Providers etc.)

» **Consulting Companies** that are providing quality engineering services to organisations of Private and Public sector.

» **Academic and Research Institutes** dealing with research on Software Engineering, DIA and Big Data technologies.

» **Public Sector organisations or NGOs** that are dealing with Big Data analysis and need to upgrade the existing infrastructure or obtain a modern one in order to enhance and enrich the quality of the provided services.

How will DICE benefit end-users?

DICE:

» Defines a quality-aware Model-Driven Architecture methodology inspired by DevOps for data-intensive cloud applications.

» Advances data-intensive-applications by offering an open source development environment.

» Provides tools for the quality assessment of Data Intensive Applications.

» Helps in customizing Big Data technologies and architectures.

» Maintains a repository of artifacts for reuse by the community.

» Uses a flexible cloud environment for deployment and production.

How can DICE improve efficiency and security?

DICE can be seen as part of the **DevOps** movement as it provides a set of tools that facilitates flow of information from Dev to Ops mediated by a model-driven approach and enables operations monitoring and anomaly detection capabilities to facilitate flow of information from Ops to Dev.

Moreover, DICE offers a number of tenets in the context of **DevOps**. Firstly, the DICE profile enables the **DIA design** with proper operational annotations that facilitate design time simulation and optimizations. 'Dev' and 'Ops' can work side by side to improve and enhance application design using simulation and optimization tools that offers refactoring of architectural designs. Secondly, DICE includes **monitoring and tools that exploit the monitoring data** to detect anomalies and to optimise the configurations of applications, also the monitoring data is exploited by tools to offer feedback to the architectural design to identify the bottlenecks and refactor architecture. Thirdly, DICE also offers tools for **automating delivery and continuous integrations** facilitate integrating design-time and runtime tools in DICE and also offering DevOps automations.

Cloud Characteristics:

❯ Low-cost software
❯ Resilient computing

HyVar – Scalable Hybrid Variability for Distributed Evolving Software Systems

Project start: February 2015
Project end: January 2018
www.hyvar-project.eu

The business potential of HyVar lies in the market need to increase maintainability and agility of distributed software. HyVar will provide small-, medium- and large-sized companies with both a methodology for adaptation of their software product lines and a software tool chain that implements it into their IDEs (Integrated Development Environment). The specific development made with regard to the HyVar architecture, tool chain and algorithms has a lot of potential in the area of highly distributed software development for connected cars, smart objects, internet-of-things and cloud-based applications.

Who is HyVar designed for?

HyVar targets **end users** who want good quality software, which responds exactly to their needs; **device developers** who need to deliver new devices as fast as possible on the market and support the reconfiguration of these devices with requirements to high distribution, flexible maintenance and quality assurance; and **software developers** who need to produce high quality software for the devices as fast as possible, responding to customer needs and to the device developers' objectives.

The pure IT solution can be exploited as software package (i.e., HyVar platform and S/W components) or as a service (i.e., customisation, maintenance, installation and training). The methodology is also offering clear advantages as it reduces the effort spent in developing and maintaining software product lines while addressing highly variable scenarios and can be exploited as consultancy services and specialized training offered to customers interested in deploying similar solutions and infrastructures.

How will HyVar benefit the end-user?

HyVar aims to support developers to **encompass unanticipated evolution** as a standard feature of software systems in production, through HyVar Domain Specific Variability Language (DSVL) and HyVar tool chain. It also exploits the **software variability** to track the exact software configurations deployed on remote devices and to enable (i) the collection of data from the devices to monitor their behaviour and perform statistical analyses; and (ii) secure and efficient highly customized software updates through HyVar cloud infrastructure. Finally, HyVar supports **over-the-air updates** of distributed applications in heterogeneous environments and enabling continuous software evolution after deployment on complex remote devices that incorporate a system of systems.

How can HyVar improve efficiency and security?

The HyVar project supports efficient software reuse integrating software product line engineering principles with existing tools and commonly used industrial practices, supporting the development and deployment of individualised software adaptations and thus realizing the concept of **Hybrid variability**. The HyVar Tool Chain also enables the definition of the underlying architecture of an organization's **software product line** (SPL) in terms of base product and variabilities and to derive a specific software product from an SPL by selecting delta modules, edit/link new required functions, automatically generate the code, validate the resulting executable against requirements. This results in an increased productivity at design and development time. After deployment, software needs to be customised for specific users and adapted to the device used as well as to the environmental conditions such as the physical location and context. The sensor data collected from a specific device triggers the automatic generation and deployment of the most appropriate software upgrade. This results in a very **cost-effective maintenance process.**

Cloud Characteristics:

- Board Network access
- Massive scale
- Measured service
- Rapid elasticity
- Service orientation
- Virtualisation

Vertical Markets:

- Engineering & manufacturing

RePhrase – REfactoring Parallel Heterogeneous Resource-Aware Applications

Project start: April 2015
Project end: March 2018
rephrase-eu.weebly.com

RePhrase produces new **software engineering tools, techniques** and **methodologies** for developing **data-intensive applications in C++**, targeting heterogeneous multicore/manycore systems that combine CPUs and GPUs into a coherent parallel platform. Data-intensive applications are one of the most important and commonly encountered classes of industrial application. Such applications are often potentially highly parallel and are a clear match to emerging heterogeneous parallel architectures. However, exploiting this potential effectively can be difficult: it is even harder to obtain good performance for parallel data-intensive applications than for compute- intensive applications, since many additional issues related to data management need to be taken into account. These include structuring the data to make it efficient to access and to process, placement/migration/replication of the data to allow fast parallel access, ensuring data consistency etc. RePhrase tackles these issues directly.

Who is RePhrase designed for?

Start-ups & microfirms, Small & medium enterprises, Open Source developers, Large companies, Technology providers, Research institutions.

How will RePhrase benefit the end-user?

(Near-) future data-intensive applications will need to consider **large-scale parallelism** as an essential part of their design and development process. RePhrase aims to dramatically **simplify** this process over the state-of-the-art using a **flexible semi-automated development approach** that will be built around emerging pattern-based parallel programming technology. Pattern-based programming enables abstraction over low-level parallelism details, including thread creation, communication, synchronisation, and scheduling; and also over data placement, access, migration and replication. This makes it ideal to address the intrinsic complexity of data-intensive applications with respect to parallelism

and data management. It will be supplemented by advanced refactoring, program analysis, testing, verification, dynamic adaptivity mechanisms and performance monitoring/measurement tools as part of a coherent development methodology. We will evaluate our work using applications taken from a number of domains, to demonstrate improvements in productivity, reliability, robustness and resilience.

How can RePhrase improve efficiency and security?

RePhrase will provide a significant **productivity increase** in the development, testing, verification, deployment and maintenance of parallel systems. It will also impact on the availability and **market take-up** of innovative tools directly, through the production of new and innovative tools for parallel systems. Finally, RePhrase will provide evidence of potential for **productivity gains** in the production of parallel, data-intensive software.

The impact of the project will be enhanced by the inclusion of a major software development company (IBM), by the involvement of 2 SMEs (Programming Research Ltd and EvoPro) and an industrially-focused research organisation (SCCH), by engagement with leading standards bodies (e.g. the C++ committee and MISRA), and by the inclusion of leading relevant European experts on parallelism and software engineering.

Vertical Markets:

- ❯ Digital health
- ❯ Energy
- ❯ Engineering & manufacturing
- ❯ Finance & insurance
- ❯ Media
- ❯ Smart cities

SUPERSEDE - SUpporting evolution and adaptation of PERsonalized Software by Exploiting contextual Data and End-user feedback

SUPERSEDE

Project start: May 2015
Project end: April 2018
www.supersede.eu

The SUPERSEDE project proposes a **feedback-driven approach for software life cycle management**, with the ultimate purpose of **improving users' quality of experience**. Decisions on software evolution and runtime adaptation will be made upon analysis of end-user feedback and large amount of data monitored from the context. An integrated platform will articulate the methods and tools produced in the project.

The project will provide advancements in several research areas however, the major contribution will be in integrating methods and tools from the mentioned areas, thus providing a new solution framework for software evolution and adaptation for data-intensive applications.

Who is SUPERSEDE designed for?

The SUPERSEDE use cases are representative of two different data-intensive application domains, **namely energy consumption management in home automation** and **entertainment event webcasting.** They address different stakeholders, including application developers, home automation platform providers, webcasting platform developers, media content designers and the end- users of software applications and services. This diversity allows a validation of the methods and tools produced to ultimately provide evidence of potential for productivity gains, and provides a wide range of opportunities for project results exploitation.

How will SUPERSEDE benefit end-users?

The SUPERSEDE project will provide advancements in several research areas, from end-user feedback and contextual data analysis, to decision making support in software evolution and adaptation. But the major novel contribution will be in integrating methods and tools from the mentioned areas, thus providing a new solution framework for software evolution and adaptation for data-intensive applications.

How can SUPERSEDE improve efficiency and security?

SUPERSEDE will deliver methods and tools to support **decision-making** in the evolution and adaptation of software services and applications by exploiting end-user feedback and runtime data, with the overall goal of improving end-users' quality of experience.

SWITCH – Software Workbench for Interactive, Time Critical and Highly self–adaptive Cloud applications

SWITCH

Project start: February 2015
Project end: January 2018
www.switchproject.eu

SWITCH addresses the urgent industrial need to develop and execute **time-critical applications** in Clouds. Applications such as disaster early warnings, collaborative communication and live event broadcasting can only realise their expected business value when they meet critical requirements in terms of performance and user experience.

Who is SWITCH designed for?

SWITCH targets:

» Software industry: to support software development and consultancy companies in delivering time-critical applications and services.

» Cloud service providers: to enable SLAs for time-critical services.

» Telecom service providers: for network providers and infrastructure operators.

» SMEs and entrepreneurs: for operating and developing their own applications with time critical requirements.

» Education organisations / Universities: for education/training purposes.

» For a wide collection of domains that require time critical services: Time critical applications in specific domains.

» Technology vendors including API management companies SDN and virtualization vendors, Telecom-managed service providers, and wireless/mobile infrastructure providers.

How will SWITCH benefit end-users?

The very high requirements posed on network and computing services, particularly for well-tuned software architecture with sophisticated data communication optimisation, implies that development of such time-critical applications is often customised to dedicated infrastructure, and that system performance is difficult to maintain when the infrastructure changes. This fatal shortcoming in the existing architecture and software tools yields very high development costs, and makes it difficult to fully utilize the virtualised, programmable services provided by networked Clouds to improve system productivity.

How can SWITCH improve efficiency and security?

SWITCH aims to improve existing **development and execution models** of time-critical applications by introducing a novel conceptual model (application-infrastructure co-programming and control model), in which application QoS/QoE, together with the programmability and controllability of the Cloud environments, can all be included in the complete lifecycle of applications. Based on this conceptual model, SWITCH provides an interactive environment to develop applications and control their execution, a real-time infrastructure planner to deploy applications in Clouds, and an autonomous system adaptation platform to monitor and adapt system behaviour.

Cloud Characteristics:

› Board Network access
› Low-cost software
› Massive scale
› Measured service
› On-demand self-service
› Rapid elasticity
› Resource pooling
› Service orientation
› Technical
› Virtualisation

Vertical Markets:

› Finance & insurance
› Media
› Smart cities

Home–grown, European Cloud and Software Services

Boosting public sector productivity and innovation through cloud computing services – Exploitable outputs ready in 2018

HNSciCloud – Helix Nebula – The Science Cloud

Project start: January 2016
Project end: June 2018
www.hnscicloud.eu

Over the coming 10-15 years, the generation of vast amounts of data created by scientific research domains will create enormous challenges for capturing, managing and processing of this data. Today, commercial cloud services do not play a significant role in the production computing environments for the publicly funded research sector in Europe. HNSciCloud will pull together commercial cloud service providers, publicly funded e-Infrastructures and the buyers' in-house resources to build a **hybrid cloud platform** on top of which a **competitive marketplace of European cloud players** can develop their own services for a wider range of users. This project will bring Europe's technical development, policy and procurement activities together to remove fragmentation and maximise exploitation.

Who is HNSciCloud designed for?

The HNSciCloud services will be made available to multiple user groups including High Energy Physics, Astronomy, Life Science, Photo/Neutron Sciences, scientists and researchers operating in the Long Tail of Science.

How will HNSciCloud benefit end-users?

The HNSciCloud project will establish a European hybrid cloud platform to support the deployment of high-performance computing and big-data capabilities for scientific research. It will address a set of challenges: compute and storage, network connectivity and federated identity management, and service payment models.

How can HNSciCloud improve efficiency and security?

The procured hybrid cloud platform will address the extreme needs of world class scientific research, providing a catalogue of secure and interoperable services from multiple suppliers; agile procurement process suitable for the dynamic cloud services market and tailored to the needs of the public research sector; development of monitoring frameworks to ensure compliance with international security and interoperability standards, performance criteria and financial benchmarking against global market leaders.

HNSciCloud will also reinforce Europe's cloud industry by reducing fragmentation across member states and permitting suppliers to offer trusted interoperable services and will improve cost-effectiveness of public sector IT systems through efficient joint procurement of cloud services (e.g. based on the adoption of standardised solutions and mutual recognition of accreditation requirements).

Cloud Computing – Exploitable outputs ready in 2018-19

ACTiCLOUD – ACTivating resource efficiency and large databases in the CLOUD

Project start: January 2017
Project end: December 2019
acticloud.eu

ACTiCLOUD proposes a novel cloud computing architecture for drastically improved management of cloud resources, targeting 1.5x increase in resource efficiency and more than 10x in scalability. Utilizing hardware intelligence that enables true resource disaggregation between multiple servers, ACTiCLOUD promotes **holistic resource management** both at the rack scale and across distributed cloud sites. Furthermore, the ACTiCLOUD novel architecture will enable the **cloudification of in-memory databases**, a core component for extremely demanding and highly critical classes of applications.

Who is ACTiCLOUD designed for?

ACTiCLOUD targets Cloud Service Providers (CSP), Cloud software vendors, Big data analysts, academics and researchers.

How will ACTiCLOUD benefit end-users?

Recent studies have shown that in most cloud data centres resource utilization ranges from 10% to 50%. ACTiCLOUD targets the two critical barriers that currently hinder true fluidity of cloud resources: the server barrier and the data centre barrier. ACTiCLOUD aims to extend resource disaggregation beyond storage, to processing cores and main memory, as well as provide a set of mechanisms and policies that will support application migrations between cloud sites that are geographically distributed. ACTiCLOUD focuses on the critical family of applications that utilise resource-demanding databases of various types with a goal to provide an environment capable of serving their high resource requests. The extended scalability and fluidity of resources provided by ACTiCLOUD, will allow the support for resource-demanding, in-memory databases in ACTiCLOUD-enabled cloud systems.

How can ACTiCLOUD improve efficiency and security?

By utilizing the ACTiCLOUD system software, system libraries and database systems, CSPs will drastically increase the **resource efficiency** of their infrastructure and make available large pools of resources to the executing applications. Thus, they will be able to support applications that today cannot be hosted on their infrastructure due to the requests being in excess of the server capacity; alternatively, they will be able to **unify their infrastructures**, as opposed to current practices where typical cloud workloads are serviced by typical low-end servers and resource hungry workloads by specialized machines typically outside the cloud. **Cloud software vendors** will improve their resource efficiency capabilities at the server, rack-scale, local cloud site and distributed cloud levels.

ACTiCLOUD will also enable a **rack-scale in-memory database** (MonetDB) and a rack-scale graph database (Neo4j), which will offer to Big Data analysts efficient (in terms of both performance and resource consumption) in-memory processing for an order of magnitude higher data volumes than provided today in typical cloud offerings. Finally, ACTiCLOUD will develop a complete set of algorithms and methodologies for **optimized resource allocation at the rack-scale** and **distributed clouds** that can be utilized by researchers to further advance findings on the field of resource-efficient computing.

Cloud Characteristics:

- Geographic distribution
- Low-cost software
- Massive scale
- On-demand self-service
- Rapid elasticity
- Resilient computing
- Resource pooling
- Virtualisation

CloudDBAppliance – European Cloud In–Memory Database Appliance with Predictable Performance for Critical Applications

CloudDBAppliance

Project start: December 2016
Project end: November 2019
clouddb.eu

The project aims at producing a European Cloud Database Appliance for providing a Database as a Service able to match the predictable performance, robustness and trustworthiness of on-premise architectures such as those based on mainframes. The project will evolve cloud architectures to enable the increase of the uptake of cloud technology by providing the robustness, trustworthiness, and performance required for applications currently considered too critical to be deployed on existing clouds.

CloudDBAppliance will deliver a cloud database appliance featuring:

» A **scalable operational database** able to process high update workloads such as those processed by banks or telcos, combined with a fast, analytical engine able to answer analytical queries in an online manner.

» A **Hadoop data lake** integrated with the operational database to cover the needs from companies on big data.

» A **cloud hardware appliance** leveraging the next generation of hardware to be produced by Bull, the main European hardware provider. This hardware is a scale-up hardware similar to the one of mainframes but with a more modern architecture.

Who is CloudDBAppliance designed for?

Our solution is designed for **data-intensive applications**. Many applications are being migrated to the cloud because of its convenience and ease of use. However, there is still a subset of applications, which are not widespread in the cloud. They are mainly data-intensive and time critical applications. Data intensive applications, if ever migrated, experience bad performance in the cloud as current cloud database infrastructure fails to satisfy high loads and does not provide predictable performance. The lack of predictable performance, in turn, leads to inability of time critical applications to satisfy

strict Service Level Agreement (SLAs) requirements imposed on them.

How will CloudDBAppliance benefit end-users?

Many of today's critical applications run on mainframes due to their high resilience and high performance. Unfortunately, there is still no equivalent to them in the cloud. This is precisely the vision of CloudDBAppliance: creating a cloud database appliance, which can replicate the performance and resilience of mainframes in cloud data centres. The goal is not only to match their performance and resilience, but also to expand their current capabilities.

How can CloudDBAppliance improve efficiency and security?

CloudDBAppliance will deliver a European appliance with a leading-edge hardware platform, the new Bullion generation equipped with an ultra-scalable operational database, LeanXcale boosted with its ultra-efficient storage engine able to scale up linearly to 1,000+ cores and integrated with ActivePivot that will provide real-time analytics over the operational data of LeanXcale.

The CloudDBAppliance outcomes will be validated through five real, industrial use application scenarios in three verticals: Finance/Banking: Real-Time Risk Analysis, ATM Optimization, Telco: Cell phone number mobility, Retail: Proximity marketing and real-time pricing.

Cloud Characteristics:

› Geographic distribution
› Massive scale
› Resource pooling
› Service orientation

Vertical Markets:

› Finance & insurance

CloudPerfect – Enabling CLoud Orchestration, Performance and Cost Efficiency Tools for QoE Enhancement and Provider Ranking

Project start: December 2016
Project end: November 2018
cloudperfect.eu

IaaS customers (such as SaaS providers) are looking for efficiency in performance that will help them stabilise the QoS offered towards their own customers in a measurable way that meets user-needs. **Public IaaS customers** also need a means of verifying that their Service Level Agreements are maintained. Finally, **IaaS customers** seek flexibility in deployment, management and migration between IaaS services. **Cloud providers** on the other hand, are aiming to enhance their Quality of Service (QoS). This can be achieved by suitably scheduling non-competing virtualized resources, thus enabling a minimization of user experience disruption and performance degradation. To enable optimised system planning and resources allocation, predictability with relation to neighbour caused overheads needs to be in place. This is often not possible since IaaS providers have little knowledge of the customer's applications.

Who is CloudPerfect designed for?

CloudPerfect is designed for IaaS customers (and consultants) and for IaaS providers. With CloudPerfect IaaS customers can easily select the provider that best fits their needs, tailored to their specific application type. Cloud providers on the other hand, can reach higher performance stability of their offer and improve customer trust.

How will CloudPerfect benefit the end-user?

CloudPerfect supports: cloud customers in finding the "perfect cloud" which fits their specific needs and cloud providers in improving trust of their existing and potential customers.

How can CloudPerfect improve efficiency and security?

CloudPerfect offers IaaS customers a toolkit which provides them:

» Performance rankings consideration to minimize time needed to review and select the appropriate provider;

» Clear and abstracted price/profit trade-offs, through reliable measurement tools/metrics, based on measurements performed against the compared services in a repetitive manner;

» Reliable monitoring and fast reporting of SLA

compliance of public IaaS services, adapting to the terms and conditions of each service;

» A generic, cloud provider independent, deployment and application management tool, to help avoid vendor lock-in, on the basis of different criteria and management logic;

» Means to understand the computational needs of their application and map it to the performance rankings.

CloudPerfect offers **cloud providers** a way to model different application categories usage patterns as well as predict the effect of certain categories groups, when scheduled on nearby resources, enabling the ability to better control the behaviour of the offered services. Furthermore, CloudPerfect applies **black box profiling** and dynamic categorization of hosted applications, in order to map customers running instances to the predefined application categories, thus giving insight to the provider, so that consolidation with minimized overheads can be performed. By using CloudPerfect IaaS providers will obtain higher consistency in performance ratings, less fluctuation on the monitored metrics and become more attractive to cloud users.

Cloud Characteristics:

- Measured service
- On-demand self-service
- Rapid elasticity
- Resource pooling
- Service orientation

Vertical Markets:

- Digital health
- Energy
- Engineering & manufacturing
- Finance & insurance
- International agencies
- Local public administrations
- Media
- National government agencies
- Research institutions
- Smart cities

COLA – Cloud Orchestration at the Level of Application

Project start: January 2017
Project end: June 2019
www.project-cola.eu

COLA targets developers that are implementing industry and/or public sector applications that require resource scalability and efficient resource utilisation. While IaaS clouds typically offer elasticity, applications cannot automatically utilise these features. Developers need to build in custom code to every single application in order to support its automatic scaling up and down. The aim of the COLA project is to ease this pain of application developers regarding **automated scalability**.

COLA develops a generic set of services (called **MiCADO**) that provide **dynamic and automated resource scalability** on clouds for all (or at least a very wide range of) applications. MiCADO services can be expressed with a set of well-defined (standardised) interfaces to be easily embedded into application control code without changing the original application logic. Moreover, MiCADO services can connect to multiple cloud middleware (e.g. EC2, CloudSigma etc.) or generic cloud access layer (e.g. CloudBroker Platform) via a set of well-defined (standardised) interfaces too. Data consistency, data protection, trustworthiness and access control are handled in a secure way when sharing and migrating applications within a single cloud or between heterogeneous cloud platforms.

Who is COLA designed for?

The **MiCADO toolkit**, the primary product of the COLA project, is specifically designed for **generic application developers and application operators**. These users are often highly skilled in relation to generic software development and systems administration skills and principles. However, they may not have deep and specific knowledge regarding cloud computing technologies.

How will COLA benefit end-users?

COLA and the MiCADO toolkit will enable generic application developers who wish to utilise cloud resources, to implement applications that automatically optimise cloud resource utilisation based on application developer/operator defined quality of service parameters. Application developers will be able to extend their application code with MiCADO service calls in order to utilise the scalability and optimisation services at both deployment and also at run-time. Application operators will be able

to define desired quality of service parameters, e.g. maximum response/completion time, maximum cost, security policy requirements etc. MiCADO services will assure that the application is optimally deployed based on the defined QoS parameters. Moreover, MiCADO services will monitor the application at run-time and will automatically scale it up or down in order to optimise application execution based on the user-defined multidimensional set of QoS parameters. Users/application operators can also modify the QoS parameters during run-time to trigger rescaling of resources, if necessary.

How can COLA improve efficiency and security?

The MiCADO toolkit will help generic application developers and operators to develop, deploy and manage cloud-aware applications. The toolkit will improve the development of cloud aware applications with no need to custom-develop every single application from scratch. By embedding MiCADO calls into the application **logic**, **scalable**, **cost effective** and **secure applications** can be **developed faster**. Additionally, MiCADO assures that the developed applications are **optimally deployed** on heterogeneous cloud computing resources based on user defined quality of service (QoS) parameters such as performance, cost and security policies. Finally, MiCADO assures that the application is managed based on the QoS parameters during run-time, and that it scales automatically to fulfil the defined **QoS requirements**.

Cloud Characteristics:

- Advanced security
- Geographic distribution
- Low-cost software
- Massive scale
- Measured service
- On-demand self-service
- Rapid elasticity
- Resilient computing
- Service orientation
- Virtualisation

Vertical Markets:

- Engineering & manufacturing

Home-grown, European Cloud and Software Services

DITAS – Data-Intensive Applications Improvement by Moving Data and Computation in Mixed Cloud/Fog Environments
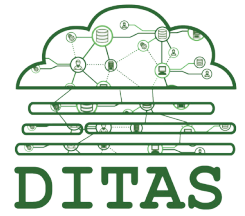
Project start: January 2017
Project end: December 2019
www.ditas-project.eu

DITAS provides real world tools targeted to simplify the life of developers. It brings to your **developer toolbox the best of Cloud & Edge worlds**.

DITAS delivers proper data management that empowers the creation of modern and decentralized data-driven applications in the post-mobile world (Cloud and Edge/Fog), delivering an Open Source SDK and a framework with tools that offer a transparent layer that takes care of all the complexity behind data management and consumption in heavily decentralized scenarios abstracting developers from details of underlying heterogeneity as well.

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, no matter the number of different devices, their type and the heterogeneity of runtime environments.

Who is DITAS designed for?

DITAS targets **application developers** that need to deploy its application in a Fog environment: Cloud+Edge and want to access data sources in the application in the same way independently if they are in the Edge or Cloud leaving DITAS to worry about moving the data between both environments. It also targets **data providers** that want to give access to its data in Cloud and Edge environments via a common API so it can be accessed by any application using the DITAS platform (Data provider and application developer could be the same entity).

How will DITAS benefit end-users?

We focus on providing real-world tools targeted to simplify the life of developers of post-mobile world applications who don't want to deal with data management details but require dealing with data in an effective, fast, agile, and secure manner. DITAS proposes a novel approach for developing data-intensive applications based on **Data Profiling tools** and **Virtual Data Containers** (VDC), 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 accurate by hiding the complex underlying infrastructure composed by different platforms, storage systems, and network capabilities.

How can DITAS improve efficiency and security?

DITAS facilitates all aspects of the data logistics in Fog environments or just pure Edge or Cloud environments (including Cloud or Edge federation). DITAS enables **application developers**, in this case data user, to develop their application using the DITAS SDK. This means they can develop the same logic to access data both in the Edge and the Cloud in the same way. They can define a data model about the data its parameters: latency, quality, etc. and the DITAS platform will accommodate the data access by moving the data from Edge to Cloud and vice versa transparently. DITAS enables **data providers** to give data sources to users in both at Edge and Cloud environments, via the definition of DITAS Virtual Data Containers. Offering data in this common API will facilitate the adoption of those data sources by third parties or, if the data provider is the same user as the application developer, it will facilitate the building of applications that access their own data.

Cloud Characteristics:

- Advanced security
- Board Network access
- Massive scale
- Measured service
- Rapid elasticity
- Resource pooling
- Service orientation

Vertical Markets:

- Digital health
- Engineering & manufacturing
- Finance & insurance
- Local public administrations

LightKone – Lightweight computation for networks at the edge

LIGHTKONE
Lightweight computation for networks at the edge

Project start: January 2017
Project end: December 2019
lightkone.eu

LightKone focusses on **moving computations** out of data centres and directly on to **edge networks**.

Who is LightKone designed for?

Software developers building services and applications on community networks, Internet of Things networks, networks of mobile devices, personal computers, and points of presence including Mobile Edge Computing.

How will LightKone benefit the end-user?

Data centres are bottlenecks for the new generation of Internet applications. For example, health monitoring applications require continuous big-data computations for each human being who is monitored. Data centres cannot scale up to support this for all human beings. Our solutions target this kind of application.

We provide solutions for **general-purpose edge computing**: we are building systems that allow to move large computations out of the data centre and directly on to edge networks. Our platforms are designed to work correctly and efficiently on edge networks with high churn and intermittent connectivity.

How can LightKone improve efficiency and security?

Running services directly on edge networks **reduces latency**, **increases scalability**, **resilience and security**, and permits **local decision making**.

Cloud Characteristics:
➤ Board Network access
➤ Low-cost software
➤ Massive scale
➤ Resilient computing

Vertical Markets:
➤ Digital health
➤ Energy
➤ Engineering & manufacturing
➤ Smart cities

MELODIC – Multi-cloud Execution-ware for Large-scale Optimized Data-Intensive Computing

Profiting from the **value** of available **data** entails running applications that process potentially large data sets, which may require significant investments in both storage and computing capabilities. Alternatively, the application can use rented infrastructures in the Cloud, although it requires mastering the offerings of one or multiple Cloud providers and increases the complexity of maintaining the running application over time. In particular if the application needs to react to external events resulting in temporary needs for more computationally expensive analysis of the data.

Not all data can be stored in with third parties in the Cloud. There are **legal constraints** on **personal data** and consumers' **privacy** concerns, and there is the **commercial data**. In addition, many providers charge if data is transferred in and out of their own Cloud than within their own infrastructure. Hence, it is a complex puzzle to find the right location for the data with sufficient computing power.

The above obstacles may prevent many companies from profiting from the value of their data or develop new businesses on commercially available or open data. In particular, it may make it difficult for Small and Medium-sized Enterprises (SMEs) to compete in the data driven economy.

Who is MELODIC designed for?

Melodic is the **Devop** for anyone who wants to benefit from the capabilities and the scalability of Cloud application execution without hassle, and at the same time be confident that the application's **data location** respects the necessary **privacy requirements** and that the data can only be **accessed** by those authorised. Melodic therefore supports a mix of private and public Cloud platforms, and execute jobs where the data is located. At the same time, the application will stay **responsive** and start or scale jobs as specified by the deployment goals of the application's owning organisation.

How will MELODIC benefit the end-user?

Melodic is a solution for **autonomic and secure cross Cloud deployment**, monitoring, and context adaptation of big data applications requiring frameworks like Apache Hadoop or Spark by constantly watching the execution, and ensuring an optimised mapping of the application's data sets and processing jobs according to your goals and requirements.

How can MELODIC help you become more efficient, more secure, faster or cost-effective?

Melodic is a tool that supports **automated deployment** of both **data** and **application jobs** processing the data based on the constraints set by the organisation owning the data and the application. It overcomes the difficult placement decision, it monitors the running application to ensure that it stays with the set constraints, and automatically adapts the application aiming to maximise the application's utility for the owning organisation in the current execution context. Hence, Melodic is a **DevOp robot that continuously tries to maximise the business value of the data**.

Melodic provides a level playing field where big companies, SMEs, and academia alike can all benefit from the cost reduction of the commercial Cloud offerings, and scale their computations when needed. If you can conceive and develop an application that extracts value from data, then Melodic will take care of running your application!

Cloud Characteristics:
- Advanced security
- Geographic distribution
- On-demand self-service
- Rapid elasticity
- Resource pooling
- Service orientation

mF2C – Towards an Open, Secure, Decentralized and Coordinated Fog-to-Cloud Management Ecosystem

Project start: January 2017
Project end: December 2019
www.mf2c-project.eu

When put together, cloud and fog computing create a new stack of resources (Fog-to-Cloud F2C), creating the need for a new, open and coordinated management ecosystem. mF2C proposes a framework to set the foundations for a novel distributed system architecture that will cover this need.

Who is mF2C designed for?

Fog computing brings cloud computing capabilities closer to the end-device and users, while enabling location-dependent resource allocation, low latency services, and extending significantly the IoT services portfolio as well as market and business opportunities in the cloud sector.

» From a **Cloud provider perspective**, mF2C creates enormous opportunities for developing and extending the service chain offering, by adding Fog in the provision of services, increasing the product/service portfolio and enabling pretty new and challenging business models.

» From a **Technology provider** perspective, mF2C will boost the adoption of IoT devices and equipment and commercial development of added value services.

» From a **Service Provider** perspective, the availability of an extended platform offers them the opportunity to develop even more sophisticated services.

How will mF2C benefit end-users?

Fog computing brings cloud computing capabilities closer to the end-device and users, while enabling location-dependent resource allocation, low latency services, and extending significantly the IoT services portfolio as well as market and business opportunities in the cloud sector. With the number of devices exponentially growing globally, new cloud and fog models are expected to emerge, paving the way for shared, collaborative, extensible mobile, volatile and dynamic compute, storage and network infrastructure.

The mF2C proposal sets the goal of designing an

open, secure, decentralized, multi-stakeholder management framework, including novel programming models, privacy and security, data storage techniques, service creation, brokerage solutions, SLA policies, and resource orchestration methods. The proposed framework is expected to set the foundations for a novel distributed system architecture, developing a proof-of-concept system and platform, to be tested and validated in real-world use cases, as envisioned by the industrial partners in the consortium with significant interest in rapid innovation in the cloud computing sector.

How can mF2C improve efficiency and security?

With the number of devices exponentially growing globally, new cloud and fog models are expected to emerge, paving the way for shared, collaborative, extensible mobile, volatile and dynamic compute, storage and network infrastructure. mF2C sets the goal of designing an open, secure, decentralized, multi-stakeholder management framework, including novel programming models, privacy and security, data storage techniques, service creation, brokerage solutions, SLA policies, and resource orchestration methods.

Cloud Characteristics:

- Advanced security
- Board Network access
- Geographic distribution
- Homogeneity
- Low-cost software
- Massive scale
- Measured service
- On-demand self-service
- Rapid elasticity
- Resource pooling
- Service orientation
- Technical
- Virtualisation

PrEstoCloud – Proactive Cloud Resources Management at the Edge for Efficient Real-Time Big Data Processing

PrEsto Cloud

Project start: January 2017
Project end: December 2019
www.prestocloud-project.eu

PrestoCloud addresses how to optimise dynamically the utilisation of cloud and edge resources, and how to proactively manage the cloud resources that also reach the extreme edge of the network.

Who is PrEstoCloud designed for?

We target **big-data solution providers** who ultilise cloud and edge resources for delivering their solutions to clients and who seek to optimise the utilisation of resources and to deliver superior user experience and improved Quality of Service.

How will PrEstoCloud benefit the end-user?

An efficient **real-time big data processing** regarding **costs and computational processing load** will be achieved by the implementation of a dynamic, distributed software architecture for proactive cloud resources management that also reaches the extreme edge of the network. PrEstoCloud will make substantial research contributions in the cloud computing and real-time data intensive applications domains, in order to provide a dynamic, distributed, self-adaptive and proactively configurable architecture for processing Big Data streams. In particular, PrEstoCloud aims to combine real-time Big Data, mobile processing and cloud computing research in a unique way that entails pro-activeness of cloud resources use and extension of the fog computing paradigm to the extreme edge of the network.

The envisioned PrEstoCloud solution is driven by the micro-services paradigm and has been structured across five different conceptual layers:

1. Meta-management;
2. Control;
3. Cloud infrastructure;
4. Cloud/Edge communication;
5. Devices, layers.

This innovative solution will address the challenge of cloud-based self-adaptive real-time Big Data processing, including **mobile stream processing** and will be demonstrated and assessed in several challenging, complementary and commercially-promising pilots. There will be three PrEstoCloud pilots from the **logistics**, **mobile journalism** and **security surveillance**, application domains. The objective is to validate the PrEstoCloud solution, prove that it is domain agnostic and demonstrate its added-value for attracting early adopters, thus initialising the exploitation process early on.

How can PrEstoCloud improve efficiency and security?

The PrEstoCloud platform will enable proactive, real-time big data, cloud-driven systems that:

» Sense the need for adapting data-intensive services proactively.

» Define on-the-fly the most suitable changes in the real-time processing architecture, including off-loading of processing tasks at the edge of the network.

» Predict reconfigurations in the underlying cloud computing infrastructure resources.

» Optimize continuously the infrastructure performance.

Cloud Characteristics:
❯ Advanced security
❯ Massive scale
❯ Rapid elasticity
❯ Service orientation

Vertical Markets:
❯ Digital health

RECAP – Reliable Capacity Provisioning and Enhanced Remediation for Distributed Cloud Applications

Project start: May 2016
Project end: December 2018
recap-project.eu

The RECAP project develops a radically novel concept in the provision of cloud services, where services are **elastically instantiated and provisioned, close to the users** that actually need them via **self-configurable** cloud computing systems.

Who is RECAP designed for?

Operators of data centers and Tier 1 hyperscale cloud service providers.

How will RECAP benefit the end-user?

The principal objective of RECAP is to reverse the current common practice to provide cloud-based services by allocating data centre resources on a best-effort basis. While recent years have seen significant advances in system instrumentation as well as data centre energy efficiency and automation, computational resources and network capacity are often provisioned using best-effort models and coarse-grained quality-of-service (QoS) mechanisms, even in state-of-the-art data centres.

These limitations are seen as a major hindrance in the face of the coming evolution of the Internet of Things (IoT) and the networked society, which are projected to significantly increase the load on networks and data centres, as well as require a much higher degree of intelligent automation.

How can RECAP improve efficiency and security?

RECAP aims to advance cloud and edge computing technology by making **application placement**, **infrastructure management**, and **capacity provisioning**, **autonomous**, **predictable** and **optimized**. It incorporates a much more elastic model, which delivers services and allocates resources in a dynamic manner, tied to time-varying user requirements.

This will ensure that **communication critical applications** will always achieve their goals without unnecessary delays, no matter where they are located. This, in turn, will minimise operational costs and improve effectiveness and energy efficiency of data centre resources.

Cloud Characteristics:

- Board Network access
- Geographic distribution
- Massive scale
- Measured service
- On-demand self-service
- Rapid elasticity
- Resilient computing
- Resource pooling
- Service orientation
- Technical
- Virtualisation

RESTASSURED – Secure Data Processing in the Cloud

Rest Assured

Project start: January 2017
Project end: December 2019
www.restassuredh2020.eu

RestAssured provides a reference architecture for **protecting sensitive data and applications** in a cloud environment by supplying secure enclaves and sticky policies for data protection that migrate across the cloud with the data. Additionally, the overall security of the end-to-end service is analysed by models@runtime and risk analysis tools.

Who is RestAssured designed for?

The RestAssured solution is a blueprint for both cloud providers, service providers and service developers.

How will RestAssured benefit the end-user?

Secure cloud computing is key for business success and end-user adoption of **public and hybrid clouds**. Currently however, most service providers would consider it highly risky to process sensitive data on public clouds, due to the vulnerability of attacks leaking data.

However, recent advantages in hardware, such as Intel's SGX (Software Guard Extensions) which provides secure enclaves for processing data, protecting applications and data against unauthorized access and tampering. RestAssured provides an architecture and reference implementation based on real-world use cases to not only securely implement cloud-based applications taking advantage of secure enclaves, but

also to provide enforcement for user/organisational defined policies for access control to sensitive data.

How can RestAssured improve efficiency and security?

RestAssured as all about the **security of data** on the cloud. Applications implemented in a RestAssured environment utilize hardware-based secure enclaves to guarantee that the code and data in the enclave cannot be externally accessed or modified.

RestAssured tools will analyse the implementation of a cloud-based service to determine if there are any security vulnerabilities. This will provide a more secure environment for both end-users and service providers on the cloud.

Cloud Characteristics:

› Advanced security

Vertical Markets:

› Digital health
› Finance & insurance
› Media
› Smart cities

UNICORN – A novel framework for multi-cloud services development, orchestration, deployment and continuous management fostering cloud technologies

Project start: January 2017
Project end: December 2019
unicorn-project.eu

UNICORN aims to simplify the **design, deployment and management of secure and elastic –by design- multi-cloud services**.

This will be achieved by:

» development and design libraries that will provide security enforcement mechanisms, data privacy restrictions, monitoring metric collection and resource management;

» enabling continuous orchestration and automatic optimization of portable and dynamic cloud services running on virtual instances or micro-execution containers for increased security, data protection privacy and vast resource (de)-allocation.

The innovation activities in the UNICORN project will be based upon existing solutions and developments to the largest possible extent, and build upon the S&T results of European RIA projects including **CELAR**, **PaaSport**, **PaaSword** and **ARCADIA**.

Who is UNICORN designed for?

Software developers in particular working for SMEs and Startups

How will UNICORN benefit the end-user?

Unicorn will enable software developers to design and develop secure and elastic applications.

How can UNICORN improve efficiency and security?

UNICORN will **reduce software release time** and provide a powerful tool for SMEs to improve software design and continuous productivity enhancement.

Software Technologies - Exploitable outputs ready in 2019

COEMS – Continuous Observation of Embedded Multicore Systems

Project start: November 2016
Project end: October 2019
www.coems.eu

COEMS will provide the first comprehensive **online observation** approach that is **non-intrusive**, allowing improved **testing** and **debugging**, and redefine the state-of-the-art for software systems development.

Who is COEMS designed for:

COEMS will provide a deeper insight to embedded systems and so improve the safety and reliability of any kind of safety-critical industries such as **automotive**, **aviation** or **medical**.

How will COEMS benefit end-users?

The main idea of COEMS is to develop and validate an efficient **real-time access** and **analysis** for operating **safe software systems**. It gives insight to the system's actual behaviour without affecting it, allowing new verification methods. Technologies developed in the project will affect:

» Software tests & Debugging.
» Performance Optimization.
» Software Engineering Research.

The COEMS approach will shorten development cycles, improves the chances of project success, lowers overall costs and is a contributing factor to a product's reliability.

How can COEMS improve efficiency and security?

Testing and debugging tasks accounts for 50 % or even more of development costs. Thus, any improvement of efficiency of developers directly translates into **increased revenue**.

Additional costs will be caused by remedying issues that occur in the field, opportunity costs of software products that are delayed and/or cancelled due to bugs (40 % of all embedded software development projects are behind schedule), consequential damage to software companies' reputations due to bugs in released software and delays and cancellations of software projects.

Cloud Characteristics:

- Advanced security
- Board Network access
- Homogeneity
- Measured service
- Resilient computing
- Resource pooling
- Virtualisation

Vertical Markets:

- Digital health
- Energy
- Engineering & manufacturing
- Finance & insurance
- Media
- Smart cities

CROSSMINER – Developer-Centric Knowledge Mining from Large Open-Source Software Repositories

Project start: January 2017
Project end: December 2019
www.crossminer.org

Developing new software systems by reusing existing open source components raises challenges in searching for open source candidate components; evaluating and selecting the most suitable open source components from amongst identified candidates; and adapting selected open source components to fit specific requirements of new software products and services.

CROSSMINER **enables in-depth analysis, evidence-based selection and monitoring of open source components**, and facilitates **knowledge extraction** from large open source software (OSS) repositories to support the development of systems for a wide range of applications by reusing the best available OSS. Selected open source components are monitored to raise alerts related to quality, and to give suggestions to reduce development effort and increase the quality of new products and services.

Who is CROSSMINER designed for?

Developers of software systems applications that wish to use OSS components for new or existing systems; providers of open source software forges; and **software developers** that contribute to the evolution of open source software projects.

How will CROSSMINER benefit the end-user?

CROSSMINER helps developers choose the right OSS by automatically extracting the required knowledge and injecting it into developers' IDE when they require it for making design decisions. This reduces time spent acquiring knowledge and improves code quality.

CROSSMINER uniquely combines advanced software project analyses with online monitoring in the IDE. The developer will be monitored to infer which information is timely, based on readily available knowledge stored earlier by a set of advanced offline deep analyses of related OSS projects.

How can CROSSMINER improve efficiency and security?

CROSSMINER's integrated platform for the development of complex software systems will (1) enable monitoring, in-depth analysis and evidence-based selection of OSS components, and (2) facilitate knowledge extraction from large open source software repositories.

OSS analysis tools extract and store actionable knowledge from a collection of OSS projects: **natural language analysis tools** to extract community communication channels metrics and bug tracking systems; **system configuration analysis tools** to provide an integrated development and operational view; **workflow-based extractors** that simplify creation of customised analysis and knowledge extraction; **cross-project analysis tools** for understanding a wide range of OSS relationships based on developer defined similarity measures; and **advanced integrated development environments** so developers can adopt the CROSSMINER knowledge base and analysis tools while receiving project input to improve developer productivity.

Cloud Characteristics:

- Board Network access
- Low-cost software
- Massive scale
- On-demand self-service
- Rapid elasticity
- Resource pooling
- Service orientation
- Technical

Vertical Markets:

- Digital health
- Energy
- Engineering & manufacturing
- Finance & insurance
- Media
- Smart cities

Home-grown, European Cloud and Software Services

DECIDE – DEvOps for trusted, portable and interoperable Multi–Cloud applications towards the Digital Single Market

Project start: December 2016
Project end: November 2019
www.decide-h2020.eu

DECIDE provides a new generation of multi-cloud services based software framework, enabling techniques and mechanisms to design, develop, and dynamically deploy multi-cloud aware applications in an ecosystem of reliable, interoperable, and legal compliant cloud services. DECIDE will provide **architectural patterns** and needed **tools** for developers and operators of multi-cloud applications (following the DevOps approach); and tools and mechanisms for **developing and operating multi-cloud applications** that can be dynamically re-adapted to be re-deployed using new and different cloud services.

Who is DECIDE designed for?

1. Developers and operators of multi-cloud native application providers.
2. Multi-cloud native application providers.
3. Users of multi-cloud native applications.

How will DECIDE benefit end-users?

DECIDE will provide the following tools:

» **DECIDE DevOps** will provide the adequate environment to design, develop, deploy and operate multi-cloud applications following the DevOps philosophy.

» **DECIDE ARCHITECT** has a set of recommended architectural patterns to be applied in order to facilitate the design, development, optimisation and deployment of a multi-cloud native application.

» **DECIDE OPTIMUS** simulates the most adequate application topology based on a set of user driven NFR and provide recommendations on the best options for the application deployment.

» **ACSmI**: The Advanced Cloud Service meta-Intermediator (ACSmI) will offer, consume and assess trusted, interoperable, and standard cloud services where to deploy the applications.

» **DECIDE ADAPT** allows the (semi-)automatic adaptation of the application and re-deployment in another configuration when certain conditions are not met.

How can DECIDE improve efficiency and security?

DECIDE tools support application developers and operators in all phases of the SDLC and Software Operation Lifecycle (SOLC) of multi-cloud native applications. The DevOps philosophy, with more agile and smaller development cycles, with faster feedback is aimed at improving productivity, from both developers and operators. At design, pre-deployment and deployment time, **DECIDE ARCHITECT** will provide architectural patterns, modelling practices, snippets of code and deployment descriptors as well as the order in which these patterns need to be applied with the main aim of increasing productivity.

DECIDE OPTIMUS aims also at increasing productivity by supporting deployment decisions of multi-cloud applications before the deployment is actually made. Candidate deployment topologies are prioritised taking into consideration the NFR entered by the application developer, which means that (s)he does not have to decide on deployment options following a trial-error approach, but rather through an optimized one, based on big data optimization algorithms. The effort of selecting, aggregating, intermediating, contracting and monitoring legal aware and accredited cloud services offerings will also be alleviated thanks to the Advanced Cloud Service meta-Intermediator (ACSmI), and the creation of the multi-cloud SLA (MCSLA) will also be easier thanks to DECIDE DevOps framework.

Finally, **DECIDE ADAPT** offers a means to monitor the fulfilment of the pre-defined application non-functional properties (NFP)s and the MCSLAs and (semi-)automatically self-adapt and re-deploy the application. This will also help to increase the productivity of operators.

Cloud Characteristics:

- Geographic distribution
- Massive scale
- Measured service
- Resilient computing
- Service orientation

Vertical Markets:

- Digital health
- Engineering & manufacturing

ElasTest – An elastic platform to ease end to end testing

ElasTest improves the efficiency, productivity, code reusability and effectiveness of the testing process of large distributed software systems. In this regard, ElasTest provides software testers the tools to instrument the software under test, orchestrate complex tests from simpler ones, and recommends good test practices.

Who is ElasTest designed for?

ElasTest is a solution specifically designed for **developers** and **testers**. By leveraging ElasTest, both groups can benefit from ElasTest in different ways.

How will ElasTest benefit the end-user?

The ElasTest solution is oriented to software developers and software testers to improve the **efficiency**, **productivity** and **code reusability** of the testing process in large complex distributed applications; and to improve the effectiveness of the testing process and, with it, the quality of the system under test.

ElasTest is specifically focussed on end-to-end and integration tests, where the application is assessed from the end-user point of view. Hence, **Quality Assurance departments** can benefit from ElasTest's capacity for running and defining complex tests from simpler ones. ElasTest provides good visualisation of these end-to-end tests, enables the definition of new tests from existing ones, and allows testers to define the conditions to be applied to the software to test it on real conditions.

How can ElasTest improve efficiency and security?

ElasTest can help you become **more efficient and faster** by reducing the time since you find a defect in your software until you find what's causing it. It helps you become more secure, because software that has been tested using ElasTest is more confident. It helps you become cost-effective by reducing the time-to-market for bug fixings, by reducing the efforts in building tests, and by enabling new kind of end-to-end tests. All these benefits led to shipping products with fewer errors.

Cloud Characteristics:

- Board Network access
- Massive scale
- Measured service
- On-demand self-service
- Rapid elasticity
- Resource pooling
- Service orientation
- Virtualisation

Vertical Markets:

- Engineering & manufacturing

OpenReq - Intelligent Recommendation Decision Technologies for Community–Driven Requirements Engineering

OpenReq

Project start: January 2017

Project end: December 2019

http://openreq.eu

OpenReq will revolutionise the quality and efficiency of **Requirements Engineering** by developing a new recommendation and decision methods inspired by advances in recommender systems, data and text mining, decision-making techniques, social networks, knowledge-based systems, semantic technologies, requirements reuse, interdependency detection, and latest approaches to user participation.

Who is OpenReq designed for?

OpenReq supports communities and single stakeholders in the **gathering and management of software requirements**. Those stakeholders include requirements analysts, project managers, software architects and developers, end-users, tool vendors, researchers, software-enabled businesses, and standardization bodies. Moreover, more unstructured bodies, like **open source communities** can benefit from the approaches proposed in the OpenReq framework.

How will OpenReq benefit end-users?

Requirements Engineering (RE) is considered one of the most critical activities in ICT projects, and poorly implemented RE is a significant risk for project failure. Requirements are the basis for all subsequent development, testing, deployment, and maintenance activities. Stakeholders typically have different backgrounds (as legal, management, technical, etc.) and thus "speak different languages" and have different priorities and expertise.

Therefore, RE is inherently **multi-disciplinary**. Moreover, due to the more and more distributed nature of software projects, regular face-to-face meetings are becoming harder to organize and stakeholders must increasingly rely on collaborative tools. Most traditional requirements management tools fail to provide adequate support for large distributed projects, even though many of them claim to do so. Multiple stakeholders can work together to define requirements, but no real support for solving conflicts, mediating tacit knowledge, or organizing stakeholders and their needs into groups exist. These are some of the reasons that makes individual as well as group decisions about the software, its features,

and services much more challenging.

How can OpenReq improve efficiency and security?

OpenReq will address requirements engineering challenges by developing an innovative, open source **recommendation and decision support platform** that assists distributed stakeholders to collaborate in requirements-related tasks. OpenReq will allow single stakeholders, groups of stakeholders, and communities to develop, discuss, negotiate, estimate, configure, and manage requirements in a context-aware and personalized way. It will also **bridge the gap** between the **development** and **usage** of **software-enabled products and services**: by taking the user community as part of the innovation process and by continuously observing and involving stakeholders and end users and tightening their commitment in the decision-making. Therefore, OpenReq will offer **support in decision making at 360 degrees**, where the end users play an important role in undertaking requirement engineering activities by providing their feedback in an explicit way-for example via social media of e-democracy platforms-or implicitly by letting the platform studying their behaviours and patterns. This will result in releasing feature that better represents the needs of the users in shorter time. In the long-term, the impact of OpenReq is to **minimize the effort** required by requirements engineering tasks by **reusing patterns and assets** produced during previous projects.

Cloud Characteristics:

- Advanced security
- Geographic distribution
- Homogeneity
- Low-cost software
- Massive scale
- Measured service
- On-demand self-service
- Rapid elasticity
- Service orientation
- Virtualisation

Q-Rapids – Quality–aware rapid software development

Project start: November 2016
Project end: October 2019
www.q-rapids.eu

How will Q-Rapids benefit the end-user?

Optimal software quality asks for the appropriate integration of quality requirements (QRs) in the software life-cycle. However, software development methodologies still provide limited support to QR management which is utterly important in rapid software development processes (RSDP): faster and more frequent release cycles should not compromise software quality.

Q-Rapids defines an **empirical-based, data-driven quality-aware rapid software development methodology**. QRs are incrementally elicited and refined based on data gathered both during development and at runtime. This data is elaborated into quality-related key indicators presented to decision makers through a strategic dashboard with advanced capabilities. Selected QRs are integrated with functional requirements for their unified treatment in the RSDP.

STAMP – Software Testing AMPlification

Project start: December 2016
Project end: November 2019
www.stamp-project.eu

STAMP provides new tools for **automatic software testing**, focusing on amplifying existing test assets. These tools automatically generate and run variants of existing unit test cases and test configurations, to detect regressions. STAMP will deliver these tools as an **Automatic Test Amplification** a **cloud service for DevOps**.

Who is STAMP designed for?

Software editors and integrators are the main users for STAMP. The test management solution market is segmented into the following three categories:

» Full-suite vendors.
» Independent pure-play vendors.
» Open-source tools.
» The DevOps community in general is the first user of STAMP.

How will STAMP benefit end-users?

The current state of practice for testing in DevOps relies on extensive manual effort to produce test cases that hardly cover all the code. STAMP's automatic amplification solutions will increase code coverage and the diversity of tested behaviour. This will reduce the number of regression bugs in production, the cost of updates and enhance trust continuous delivery. This potential for radical innovations in DevOps relies on one key concept: deliver test amplification software services, which can increase test automation at multiple development steps that span from early development to operations in production and which can be integrated in various DevOps toolchains.

STAMP addresses test automation at **unit level**, **configuration level** and **production stage**, leveraging the human knowledge and the manual effort invested in the production of test assets (unit tests, APIs, manually defined test configurations, etc.) to increase their value through automatic amplification.

How can STAMP improve efficiency and security?

STAMP test amplification techniques aim to **reduce the accumulation of technical debt** due to prohibitive refactoring costs, by significantly lowering the risk of opportunistic refactoring. The impact will be a decrease in associated **long term maintenance costs** and will:

» Increase the diversity of execution paths covered by 40%.
» Decrease by 20% the number of tests which fail once but not again if run several times.
» Increase by 20% the number of lines of product code, which are executed for each second of time spent running tests.
» Increase by 40% the number of unique invocation traces between services in a global perspective.
» Increase by 30% the number of valid bugs detected during testing which are specific to the generated configurations.
» Reduce by 30% the time on configuring and deploying products for testing purpose.
» Reduce the size of log files by an order of magnitude, keeping all essential information.
» Increase by 70% the number of crash replicating test cases.
» Enhance existing test suites with 10% of production-level test cases.
» three test amplification services integrated in two different toolchains.

Cloud Characteristics:

> Advanced security
> Board Network access
> Homogeneity
> Measured service
> On-demand self-service
> Service orientation
> Virtualisation

Vertical Markets:

> Digital health
> Finance & insurance
> Media
> Smart cities

EU-Japan Research and Development Cooperation in Net Futures –
Exploitable outputs ready in 2017

iKaaS – intelligent Knowledge-as-a-Service Platform

Project start: October 2014
Project end: September 2017
ikaas.com

iKaaS is paving the way **towards trouble-free multi-cloud based services**. We envision that in this scenario, many parties can be a provider as well as a consumer of cloud-based assets, helping to extend cloud capabilities and leverage at the same time the wealth of knowledge and services empowered through this. Entities such as Public Administrations, Health Support Service Providers, Town Management Service Providers, these days have access to a wealth of data and control capabilities offered by sensor installations and other IoT devices. However, what is sometimes missing are the necessary tools to interpret and understand what the data mean, what insights they can reveal, what actions would need to be taken in response. Also, additional useful data may not be available due to technological as well as privacy related reasons.

Who is iKaaS designed for?

iKaaS will deliver a **toolbox** that can be used by **Public Administrations**, **Health Support Service Providers**, **Town Management Service Providers** as well as **Technology/Solutions Providers**. The former can benefit by using the iKaaS toolbox in order to roll-out cloud-based services; the latter can benefit by re-using or even extending iKaaS tools to be used in similar or even different contexts and enhance their product portfolio. In addition, everyday citizens with interest in technology and with suitable home infrastructure can benefit, using iKaaS tools to setup their own home cloud installations, which may then be used for multi-cloud service offerings.

How will iKaas benefit the end-user?

Thanks to iKaaS, our city council now can **exploit the information** coming from **sensors** all around the city. We can now promptly **predict situations** that can lead to **hazards** and **lost time** for them (CO2 emissions, pollen levels, floods, icy conditions, congestion). At the same time, we realised that we could do so much more with our current infrastructure by managing data and services intelligently also for additional urban services. We even **reduced our operational costs** and we were able to pass this reduction down to our citizens. Our **application developer business is growing**. We can provide better support for our customers and can handle sensitive personal data

much easier, as well as reduce the service response time. In addition, we can introduce new features to our products much faster re-using existing software resources, simply combining them in a different way. This together with a much more efficient use of cloud resources has **given us an edge** over competing solutions and has allowed us to **increase our customer base** considerably.

How can iKaaS improve efficiency and security?

The above-mentioned Service Provider stakeholders can roll-out services that leverage the extended capabilities that a multi-cloud platform brings in. Additional **privacy preserving data collection** and **control capabilities** through the IoT endpoints, additional hardware resources together with the iKaaS analytics and service management tools can help **reduce both CAPEX and OPEX costs**. This decrease in CAPEX, OPEX as well as software development costs can also lower the barriers for SMEs to enter the market. In addition to improving the quality of services that can be offered to citizens/customers, these can eventually come at a lower cost for the citizen/customer.

Cloud Characteristics:

- Advanced security
- Board Network access
- Geographic distribution
- Homogeneity
- Low-cost software
- Massive scale
- Measured service
- On-demand self-service
- Rapid elasticity
- Resilient computing
- Resource pooling
- Service orientation
- Virtualisation

Vertical Markets:

- Digital health
- Local public administrations
- Smart cities

EU-Japan Joint Call – Exploitable outputs ready in 2017-19

BigClouT – Big data meeting Cloud and IoT for empowering the citizen clout in smart cities

Project start: July 2016
Project end: June 2019
bigclout.eu

BigClouT project makes use of today's three key technologic enablers, Internet of Things (IoT), cloud computing and big data, for the objective of increasing the efficiency in using urban infrastructure, economic and natural resources shared by the increasing population. BigClouT will offer an **analytic mind** to the city by creating **distributed intelligence** that can be implanted throughout the **whole city network**.

Who is BigClouT designed for?

BigClouT provides necessary platforms and tools to help cities **manage their resources** in a more efficient manner; thus, improving the **quality of life** and **security** of its citizens.

How will BigClouT benefit the end-user?

The BigClouT project targets in particular the following three technical challenges:

- » **Interoperability**: Many IoT technologies are emerging which promise to make the IoT ubiquitous in the smart city context. However, a "universal glue" to bind the "things" to each other, to applications and to the cloud is missing. BigClouT provides an integrated data collection and redistribution platform that can universally handle data from various city information sources.
- » **Self-awareness and dependability**: Several tens of billions of devices are expected to be deployed in the coming decades. Naturally, the complexity of systems continues to grow and human administrators cannot handle this complexity with the traditional tools. BigClouT proposes mechanisms to make smart city platforms self-aware of their physical and virtual contexts and manage themselves with certain autonomy.
- » **Making value of big data**: The intelligent use of data revolutionizes decision making in businesses, sciences and society in the future, by offering predictive and prescriptive analytics. BigClouT proposes storing and processing close to the edge level, thus creating real-time distributed intelligence, which also natively facilitates privacy protection and scalability, comparing to centralized cloud solutions.

How can BigClouT improve efficiency and security?

The BigClouT project will validate its approach with real-life deployments in European and Japanese cities:

- » **Grenoble (FR):** Monitoring the impact of big international events (scientific or business) on the local economy in order to optimise the investment of the Grenoble-Alpes Métropole. Monitoring the industrial parks and propose services for the employees of the park to increase their quality of experience (e.g., transportation, information, recommendations of activities, personalized offers).
- » **Fujisawa (JP):** Improve hospitality and regional economy synergies through increading the number of international tourists through events such as Tokyo Olympics Paralympics 2020.
- » **Tsukuba (JP):** Environment and Congestion Prediction in particular for disaster prevention and for boosting local tourism economy. Providing personalized recommendations to foreign business and sightseeing tourists.
- » **Bristol (UK):** Improving transport options to maintain safety, traffic flow and an environmentally friendly urban environment. Personalised energy optimisation for homes.

Cloud Characteristics:

- Advanced security
- Board Network access
- Geographic distribution
- Homogeneity
- Massive scale
- Measured service
- On-demand self-service
- Rapid elasticity
- Resilient computing
- Resource pooling
- Service orientation
- Technical
- Virtualisation

Vertical Markets:

- Smart cities

EU-Brazil Research and Development Cooperation in Advanced Cyber Infrastructure – Exploitable outputs ready in 2018

EUBraBIGSEA – EUrope–BRAzil Collaboration on BIG Data Scientific REsearch through Cloud-Centric Applications

Project start: January 2016
Project end: December 2017
www.eubra-bisea.eu

EUBra-BIGSEA will provide an abstract framework for the development of **distributed Big Data applications**. Multiple data models are supported (e.g. data streams, multidimensional data, etc.) and efficient mechanisms ensure privacy and security, on top of a QoS-aware layer for the smart and rapid provisioning of resources in a cloud-based environment. The **Europe – Brazil cooperation** and the careful selection of the joint consortia on this project will leverage the experience gained on consolidated, past international projects in the area of cloud and Big Data services.

Who is EUBra-BIGSEA designed for?

Different project outputs target different stakeholders:

» **Architecture & components user communities**:
  » Data scientists & Domain researchers.
  » Administrators & Cloud infrastructure providers.
  » Open source communities & application developers.
  » Private sector: SMEs, Start-Ups, large enterprises.
» Use case / Application field user communities:
  » **Public sector**: Public authorities; Urban planners.
  » **General public**, citizens.

How will EUBra-BIGSEA benefit end-users?

EUBra-BIGSEA is developing a QoS architecture that will **predict resource consumption of Big Data Analytics applications** in order to **pre-allocate** and **dynamically adjust virtualized infrastructures**. EUBra-BIGSEA will leverage mixed horizontal and vertical elasticity on hybrid container and VM infrastructures to support a rich Data Analytic framework powered by OPHIDIA, COMPSs and Spark, which will be enriched with these capabilities.

How can EUBra-BIGSEA improve efficiency and security?

EUBra-BIGSEA provides an **integrated, elastic and dynamic big data cloud platform** to address knowledge discovery by tackling data volume, variety, velocity and veracity issues as well as privacy, security and QoS challenges. It is also developing a programming layer for Big Data to **transparently build applications** composed of data operators mapped to different Big Data frameworks.

EUBra-BIGSEA is also developing a framework for ensuring the **Quality of Service** of data analytics services on top of cloud computing infrastructures. It also provides a comprehensive and effective **security approach** and a global security solution that takes into account the requirements and constraints of the cloud for Big Data processing.

Cloud Characteristics:

- Advanced security
- Board Network access
- Geographic distribution
- Homogeneity
- Low-cost software
- Massive scale
- Measured service
- On-demand self-service
- Rapid elasticity
- Resilient computing
- Resource pooling
- Service orientation
- Technical
- Virtualisation

Vertical Markets:

- Smart cities

EUBrasilCloudFORUM – Fostering an International Dialogue between Europe and Brazil

Project start: January 2016
Project end: December 2017
eubrasilcloudforum.eu

Cloud computing is a catalyst for economic growth and new business opportunities in both Europe and Brazil. An **open exchange** with stakeholders from both **Europe** and **Brazil** is critical for building common consensus and for delivering practical guides on how Cloud services can help business and research activities. It is important to foresee future challenges, gaps and opportunities to get the most out of the Cloud and to keep Europe and Brazil on the cutting edge of ICT, maximising the growth potential of a transatlantic digital economy.

EUBrasilCloudFORUM enables Europe and Brazil to formulate and develop a **common strategy** and approach for Research & Innovation (R&I) in Cloud Computing in line with the priorities of each region. The project plays an instrumental role by coordinating groups of experts investigating Cloud Computing topics for **future research cooperation in ICT**, particularly cloud computing, big data and the Internet of Things (IoT).

Who is EUBrasilCloudFORUM designed for?

EUBrasilCloudFORUM supports European and Brazilian:

» **Policy makers and ICT experts** in developing a common strategy and approach for R&I and cooperation in Cloud Computing;

» **Startups and SMEs** by giving visibility to their services with both European and Brazilian audiences through **EUBrasilCloudFORUM marketplace**, and facilitating their market uptake and transatlantic business cooperation;

» **Academics and Researchers** by leveraging on past Research and Innovation initiatives and guiding current ones on the relevant topics for EU-Brazilian cooperation in ICT.

How will EUBrasilCloudFORUM benefit end-users?

The **EU-Brazil Research & Innovation Roadmap for 2018-2019**, proposes specific cloud computing topics for consideration in the context of the near future EU-BR collaborative calls. An **online marketplace**, provides a catalogue of relevant assets and players to support public and private stakeholders interested in transatlantic collaboration.

EUBrasilCloudFORUM also drives a **community driven engagement forum** for EU-Brazil policy, research & innovation dialogues engaging experts and stakeholders into two Working Groups. This is supported by a series of **EU-BR Cloud Computing events** including Cloudscape Brazil (5-6 July 2017, Sao Paolo, Brazil) and Cloud Network workshops. The project also contributes to the organisation of EU-Brazil Policy Dialogues in ICT.

How can EUBrasilCloudFORUM improve efficiency and security?

The expected impact of EUBrasilCloudFORUM is to nurture an environment to enable market uptake of EU-BR research and innovation outputs, by organizations of all sizes, types and across sector, to support the growth of a transatlantic data economy thanks to Cloud Computing. The project aims to promote cloud computing **cooperation** and **dialogue** between **Policy support** via programmed management, researchers and innovators from Europe and Brazil. Furthermore, EUBrasilCloudFORUM intends to support the formulation and development of an **EU-BR strategy in Cloud Computing** for the forthcoming years.

SecureCloud – Secure Big Data Processing in Untrusted Clouds

Project start: January 2016
Project end: December 2018
www.securecloudproject.eu

SecureCloud addresses the **confidentiality**, **integrity** and **availability** of applications executed in the cloud. The main problem that we face is how to ensure the confidentiality of data while being processed. Our approach is based on upcoming hardware extensions of commodity CPUs.

Who is the SecureCloud designed for?

Our end users are those who need to **process massive amounts of data** in a timely and secure fashion. Data at rest or in transit on the network is typically protected by **encryption**. However, when data is to be processed, it must be **decrypted**. From a security point of view, this is an issue, since privileged users and/or software might get access to the data in plain text. The project focuses on the **power grid** domain, but the features it delivers can be used in a variety of other domains.

How will SecureCloud benefit the end-user?

SecureCloud focusses on a particularly important domain: applications that support **critical infrastructures**. The security guarantees of current cloud offerings are not sufficient for such applications (as well as for a large fraction of existing and emerging applications). Hardware extensions of commodity CPUs provide powerful mechanisms that effectively protect data from attacks, even from privileged users/software. Regrettably, using CPU hardware extensions requires advanced programming skills, which a large fraction of application developers lacks. SecureCloud makes CPU extensions readily available via a **pre-packaged container**. By doing so, SecureCloud enables developers to deploy their applications in a secure way and at no extra cost.

How can SecureCloud improve efficiency and security?

By removing technical impediments to dependable cloud computing, SecureCloud will encourage and enable a **greater uptake of cost-effective, environment-friendly**, and **innovative cloud solutions**, in particular for critical infrastructure applications throughout Europe and Brazil. If one can trust a cloud to run applications in the context of critical infrastructures, one can clearly trust this cloud to run applications in a large variety of application domains. In stark contrast to traditional throughput-oriented, batch-processing cloud applications, applications in the critical infrastructure domain do not only have strong requirements with respect to confidentiality, integrity, availability, but they typically are also latency sensitive.

Cloud Characteristics

- Advanced security
- Board Network access
- Geographic distribution
- Low-cost software
- Massive scale
- Measured service
- Rapid elasticity
- Resilient computing
- Resource pooling
- Service orientation
- Virtualisation

Vertical Markets:

- Energy

Notes

CloudWATCH2 Consortium

HORIZON 2020