

CloudscapeVI

Cloud for a smart economy & smart society

24 & 25 February 2014
Brussels, Belgium

Position Papers

Sponsors

Premium sponsor



Demo-zone sponsor



Supporters

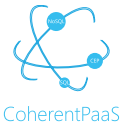
































Disclaimer

The views expressed in the position papers in this document are those of the authors and do not necessarily reflect the view of their organisations and/or affiliates. © Copyright Cloudscape Series.
Editors: Stephanie Parker, Silvana Muscella & Nicholas Ferguson, Trust-IT Services Ltd

Contents

Welcome Message 3

Perspectives from our Premium Sponsor 5

Windows Azure for Research 6

Perspectives from our Demo Zone Sponsors 8

Measuring the impact and business value of cloud computing 9

Greening with Cloud: How to gain cost savings; deliver better citizen services and provide a low carbon economy 10

The transformation impacts of the cloud model in enterprises 12

The Business Value of the Cloud and IT 15

Challenges and future directions 18

Cloud-friendly laws? 19

e-Infrastructure commons marketplace 22

The Future of Cloud in Europe – Successful adoption requires trusted clouds and BIG data 25

IEEE Intercloud Project – P2302 Working Group, and Global Testbed 27

Expert insights on data regulations, contracts and security 31

Contractual and data protection aspects of cloud computing 32

Service Level Agreements for data protection and data security 34

Cloud computing in the public sector – A European Perspective 36

Negotiating cloud contracts – from both sides now 38

Smart services, smart society – Perspectives from the public sector 40

EDM an Austrian eGovernment CLOUD Services that protects the environment and makes complex processes manageable – EuroCloud 2013 Best Cloud Service Use Case Public Sector 41

Cloud for Europe – Challenging the European market for public administrations 43

Me and My Cloud 45

Interoperability is the key to freedom in the Cloud 46

The evolution of the ~okeanos IaaS cloud service 50

Enabling Swiss Researchers with Cloud 52
Scientific cloud computing using e-Science Central 54
EMBL-EBI's Embassy Cloud: Bringing computation to large data sets 56
EU Brazil Cloud Connect – Addressing societal challenges in the cloud 58
Boosting business innovation in the Cloud 61
Flexibility in financial services through the cloud 62
Startups and small businesses in the cloud – Experiences from Cloud Software Finland 64
MobiCloud - a novel cloud-based platform for cross-platform context-aware enterprise mobile apps 66
Transmetrics – A Cloud solution that brings big data for cargo transport 68
Scalable Data Analytics – A new start-up in the cloud 70
ClouDesire – a new cloud based app store for software vendors 72
The Next Wave of European Innovation 74
CoherentPaaS: Coherent and Rich PaaS with a Common Programming Model 75
StackSync: open source personal cloud for organisations 77
SyncFree: Large-scale computation without synchronisation 80

Welcome Message

This year's Cloudscape is the 6th in a series of annual gatherings and the 2nd as a self-sustained event lead by Trust-IT Services & a prestigious Programme Committee. Cloudscape is the lasting legacy of 2 previously EC Coordination & Support Actions funded under the e-Infrastructures unit of the EC Framework Programme 7. Riding the wave of the emergence of cloud computing, the event has grown since 2009 in terms of its reputation and visibility across an increasingly diverse range of stakeholders. This year's edition zooms in on the strategic role of cloud adoption and its socio-economic benefits with practical "how-to" sessions on making the transition and overcoming some of the barriers.

We are honoured to welcome some of the most important, international thought leaders in the cloud space to offer their inspirational perspectives on where we're heading. A key driver behind Cloudscape VI is advancing common knowledge on cloud computing, taking stock of where we stand today, and examining issues like data protection regulation which affect all us. Experts will explore ways of moving towards the European goal of a trusted and secure cloud environment, and how we may remove barriers to wider uptake and measure impact. Over the next two days, you will witness views and visions from cloud providers experiences from European businesses, from large corporations to new start-ups, as well as insights from research & education, government & public authorities. Cloudscape VI showcases, for the first time, a set of small business success stories with practical experiences and guides that can help others plan their journey to the cloud. New priorities in the business world include establishing leadership with trusted cloud infrastructures and services in Europe and cloud with big data analytics.

There's some exciting stuff happening in research with cloud now enabling researchers from disciplines that have never had access to supercomputers and advanced technologies. Big science, with its big data, are also benefitting from the cloud with examples from CERN Openlab, the European Space Agency and the European Bioinformatics Institute. Moving forward, we need to make concerted efforts to achieve the e-infrastructure commons vision aimed at providing public and private researchers with access to worldwide and world-class resources and services through a dynamic and sustainable marketplace.

Cloudscape is delighted to host several side sessions, giving you extra choices, running parallel to the main programme which are open to all event participants. These sessions focus on lessons learned in cloud adoption, tackling cloud issues and uptake from a practical perspective, from contractual and data protection issues, security and compliance, interoperability and a deep dive on Windows Azure.

We also welcome back the Cloudscape Demo-zone, this year sponsored by IEEE Cloud Computing with demos running during coffee and lunch breaks on both days.

You will have plenty of opportunities to meet old friends and make new ones during the networking breaks. The Friends of Cloudscape Cocktail starts at 6pm on Monday 24 February hosted by our Premium Sponsor, Microsoft. This relaxing setting is perfect for catching up or making new acquaintances.

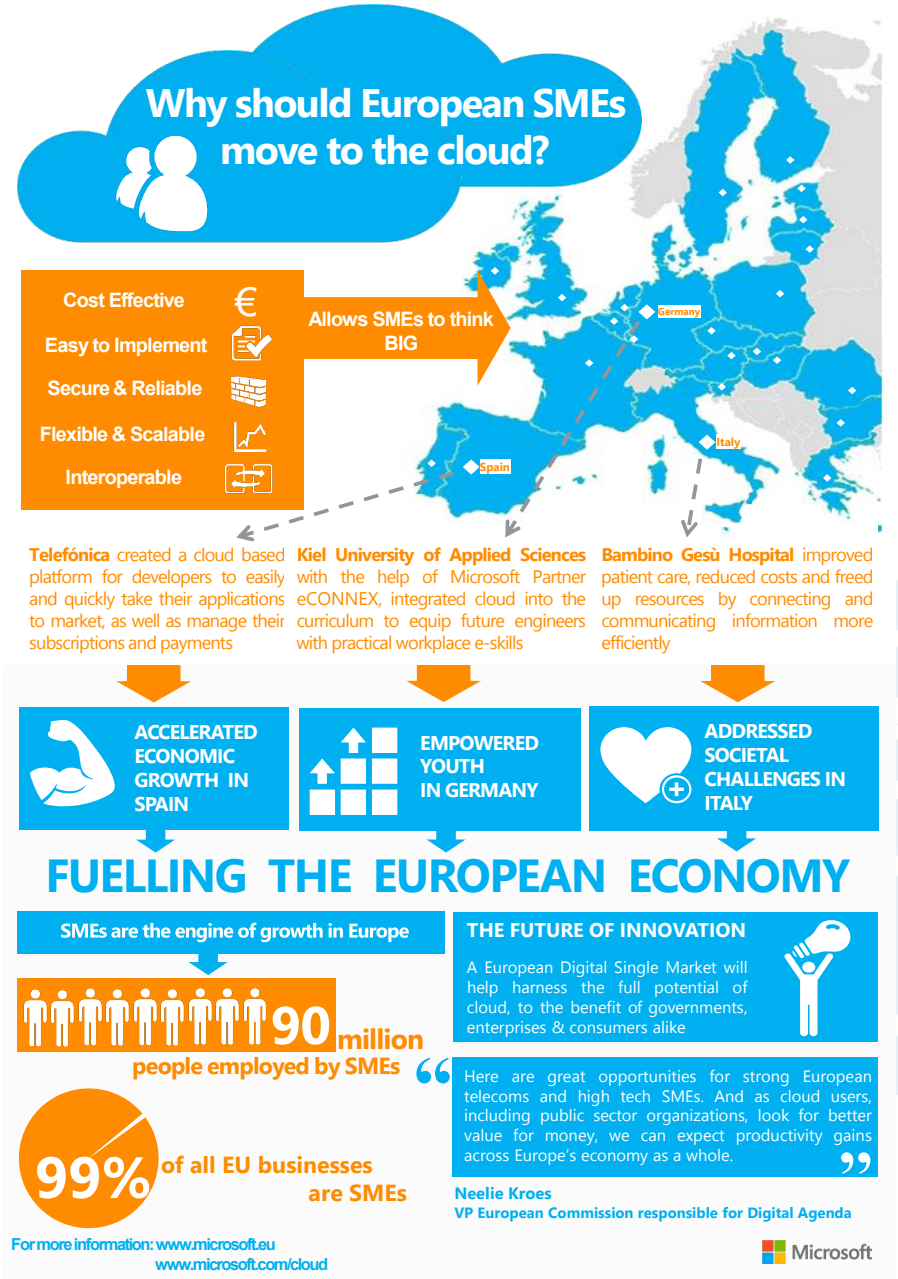
EC representatives will be wrapping up Cloudscape VI with insights on funding opportunities under Horizon 2020 on the key themes running through the 2 day packed event: Cloud Computing, Data Infrastructures, Sustainable Energy and the Future Internet, so you can travel home with new ideas already forming.

We would like to extend our thanks and appreciation to our sponsors: Microsoft and IEEE Cloud Computing, our distinguished Programme Committee presented below, all our supporters, chairs, speakers, panellists, EC representatives and, of course, all of you here today for ensuring Cloudscape advances common knowledge on the cloud.

Programme Committee

Joe Baguley, Chief Technology Officer, EMEA VMware
David Bernstein, Managing Director, Cloud Strategy Partners, IEEE Cloud Computing & IEEE Intercloud Project
Francisco Brasileiro, Assistant Professor, Computing and Systems Department, Federal University of Campina Grande
Gabriella Cattaneo, Associate Vice President, IDC European Government Consulting
Michel Drescher, Technical Manager, EGI.eu
Dean Flanders, Head of Informatics, Friedrich Miescher Institute & President, Swiss National Grid Association
Vangelis Floros, Project Manager, GRNET
Fabrizio Gagliardi, Independent Consultant, Chair of ACM Europe
Tobias Höllwarth, Chairman, EuroCloud Austria
Geir Horn, Head of European ICT projects at the University of Oslo
Gershon Janssen, Independent Consultant & Board of Directors, OASIS
Peter Kunszt, Director at the Service and Support for Science IT, University of Zurich
Steven Newhouse, Head of Technical Services, European Bioinformatics Institute
Alan Sill, Texas Tech University High Performance Computing Centre & VP Standards, OGF
David Wallom, Associate Director – Innovation, Oxford eResearch Centre

Perspectives from our Premium Sponsor



Windows Azure for Research

Dr Kenji Takeda, Microsoft Research

Advancements in technology are transforming the way we conduct research. Scientific instruments, environmental sensors, and large-scale simulations are generating more scientific data than ever before. To take full advantage of these data, researchers need powerful computing and massive storage resources. Windows Azure, Microsoft's flexible and scalable cloud computing platform, brings unlimited possibilities for affordable data-intensive computing, as well as data storage and sharing. The Windows Azure for Research program provides academic, government, and industry researchers with access to a variety of tools and resources to help them maximize cloud computing for their research [1].

Who stands to benefit and how

Researchers and e-infrastructure providers are benefiting from using Windows Azure to build and provide scalable, on-demand, easy-to-use compute, storage, data distribution and mobile services. By reducing cycle times and data access for science and research, it benefits both researcher and the general public. Cloud computing can enable better access to science and research, delivering information and services in more usable and reliable ways. This ranges across disciplines, including engineering, environmental science, the humanities and life sciences.

Our Windows Azure for Research programme aims to provide support in the form of Windows Azure Awards for cloud computing and storage, no-cost training classes across Europe and the world, webinars, support and technical resources.

Addressing new challenges

The opportunities for exploiting cloud computing to solve the challenges around open and reproducible science at scale are significant. It is clear that cloud computing is becoming a key part of the research ecosystem, and our programme is helping the community to develop new ways of working to maximise the benefits to researchers. To provide real-world impact, we are now supporting around 100 Windows Azure for Research projects globally across many domains and technology areas [2]. These projects use a variety of open platforms, including Linux, Python, and Java, with several requiring interoperability between different platforms and frameworks. In Europe, a selection of projects includes:

» *ApiSwarm: Elastic processing of crowd-based datasets*, **Romain Rouvoy, University Lille 1, France**

» *Minority Report: Using the Cloud to Enable Proactive Digital Forensic Investigations*, **Liliana Pasquale, University of Limerick, Ireland**

» *sAfe CitiEs through clouD and Internet-of-Things (ACED-IoT)*, **Elisabetta Di Nitto**, **Politecnico di Milano**, **Italy**

» *Towards an interactive secondary analysis of RNA sequencing data service in Windows Azure cloud with*

Apache Spark framework, **Marek Stanislaw Wiewiorka, Warsaw University of Technology, Poland**

» *Alzheimer Bio Project*, **Evgeny Rogaev**, **Vavilov Institute of General Genetics**, **Russian Academy of Science**, **Russia**

» *Recomputation of Scientific Experiments*, Ian Gent, University of St Andrews, United Kingdom

» *ADAM+ - A Large-Scale Distributed Image and Video Retrieval System*, **Heiko Schuldt, University of Basel, Switzerland**

» *Analysis and Interpretation of Human Exome Sequencing for Clinical Diagnosis and Electronic Healthcare Record Integration in the Cloud*, **Paolo Missier, Newcastle University, United Kingdom**

» *JASMIN: e-infrastructure for Climate and Earth System Science*, **Philip Kershaw**, **STFC Rutherford Appleton Laboratory**, **United Kingdom**

» *Real-time Catastrophe Risk Management on Windows Azure*, **Blesson Varghese, University of St Andrews, United Kingdom**

» *British Library Labs in the Cloud*, **Adam Farquhar**, **The British Library**, **United Kingdom**

This breadth and depth of activity shows how cloud computing can be applied effectively, and is helping to map the key design patterns and practices that can be reproduced by others in the research community in a practical way. Through these projects, researchers are exploring new business models to support globally scalable compute and data services and workloads using the public cloud.

A key challenge in scientific computing is that of skills and expertise. Cloud computing offers the advantage of being mainstream and is designed to be highly usable, with readily available training and technical resources. Our Windows Azure for Research programme builds on this, adding focussed efforts aimed at addressing skills and training for researchers and e-infrastructure providers directly through:

» **Hands-on training workshops** across Europe and the world, at no-cost [3]. Our materials are released under an open license, so the research and education communities can re-use and extend them for their own purposes.

» **Online training and webinars**, to provide on-demand and interactive resources [3].

» **Technical white papers** and walkthroughs targeted at researchers [4].

Our experience across the globe is showing that Windows Azure can be effectively applied across many domains to become a core part of research e-infrastructure as we move towards a *"Science 2.0"* world.

Links and References

[1] www.azure4research.com.

[2] http://research.microsoft.com/en-us/projects/azure/projects.aspx.

[3] http://research.microsoft.com/en-us/projects/azure/training.aspx.

[4] http://research.microsoft.com/en-us/projects/azure/technical-papers.aspx.

Perspectives from our Demo Zone Sponsors

IEEE CLOUD COMPUTING

Get Involved

with the IEEE Cloud Computing Initiative



Cloud Computing has widespread impact across how we access today's applications, resources, and data. The IEEE Cloud Computing Initiative (CCI) intends to lead the way by collaborating across the interested IEEE societies and groups for a well-coordinated and cohesive plan in the areas of big data, conferences, education, publications, standards, intercloud testbed, and dedicated web portal.

Get involved

The CCI offers many opportunities to participate, influence, and contribute to this technology.

Email: cloudcomputing@ieee.org

Portal: http://cloudcomputing.ieee.org

IEEE

Save the Date

- European Cloud Congress, 21-25 July 2014, Vasteras, Sweden (compsac.cs.iastate.edu), co-located with COMPSAC
- North America Cloud Congress, 8-12 December 2014, Austin, TX, US (www.ieee-globecom.org), co-located with GLOBECOM
- Cloud Computing for Emerging Markets (CCEM), 15-17 October 2014, Bangalore, India (cloudcomputing.ieee.org/ccem)
- Asia Pacific Cloud Congress, 15-18 December 2014, Singapore, (2014.cloudcom.org), co-located with CloudCom

Follow us on

twitter @ieeecloud

IEEECloudComputing

IEEE Cloud Computing

Measuring the impact and business value of cloud computing

Greening with Cloud: How to gain cost savings; deliver better citizen services and provide a low carbon economy

Robert B. Bohn, National Institute of Standards and Technology (NIST)

The NIST Cloud Computing Program was initiated in May 2010 and charged with determining the high priority requirements in security, portability and interoperability for the successful adoption of cloud computing for the US Government. As part of this effort, a cloud computing reference architecture and taxonomy were produced and generated requirements for precise language for service level agreements and cloud computing service metrics. The work supports a shared services model for delivery of better citizen services with cost savings.

Who stands to benefit and how

In cloud computing, advanced software and networking technologies provide multiple environments for multiple users in a single hardware environment. It is the sharing of compute environments which provide a reduction of physical hardware overhead, thus leading to overall savings in energy consumption. The consumers of cloud services and their customers will benefit from learning that this sharing can also be applied to a shared services model for delivery to customers as a viable method for lowering their overall energy consumption. Cloud Providers will also benefit from discussion as they will need to innovate in order to increase their overall energy efficiencies.

Sustainable services, efficiencies and impact

Cloud Computing has inherent properties which allow it to be a good value for a consumer interested in cost savings and efficiency. These efficiencies can take on many different qualities. For example, one can consider efficient delivery or consumption of services, efficient methods of processing or how one can be efficient in a multi-tenant cloud environment. These are all examples of resource efficiencies. However, at their core, efficiencies have to do with processing, storage or transmission of data which ultimately boils down to a discussion of energy efficiency.

The largest example of energy savings for a consumer comes from reducing their overall number of datacentres. In this case, they do not have the expense of power, heat and cooling for the compute systems as it is all under the provider's roof. Additionally, the move to consolidation saves money by not owning and not having to house the infrastructure which also implies an energy savings since you no longer have the expense for the lease of the building and the associated utilities.

Cloud technology also allows for other kinds efficiencies as well. Like many things in the discussion of cloud computing, the answer is found in the ability to scale up. Virtualisation allows multiple system images to be

run on a single server. There is no need to power up another server since less physical equipment is required to perform the same set of operations. The software that makes a cloud a cloud and gives it the power to rapidly provision and scale workloads will also maximise the use of the server.

A final method in which these efficiencies can be translated into operational success is through a shared services model for delivery to stakeholders and customers. This sharing mechanism builds upon the inherent nature of sharing, elasticity and provisioning in the cloud computing model.



The transformation impacts of the cloud model in enterprises

Gabriella Cattaneo, David Bradshaw, IDC EMEA

International Data Corporation (IDC) [1] is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications, and consumer technology markets, with more than 1,000 analysts and coverage of over 110 countries worldwide. IDC helps IT professionals, business executives, and the investment community to make fact-based decisions on technology purchases and business strategy. In 2014, IDC celebrates its 50th anniversary. Cloud computing is one of the key focus areas of IDC's research. IDC constantly analyses the Cloud Computing market, publishing off-the shelf and ad-hoc studies and reports on market trends, supply-demand dynamics, market size and innovation trends.

Impact of the cloud model in enterprises

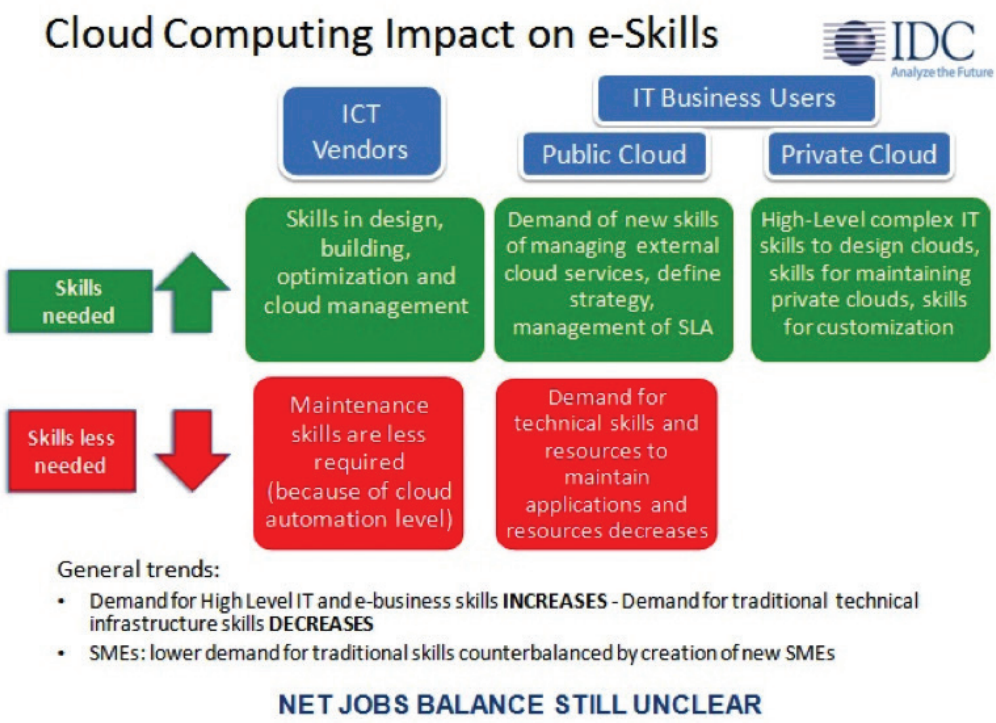
According to IDC estimates, worldwide spending on public IT cloud services reached $47.4 billion in 2013 and is expected to be more than $107 billion in 2017. Over the 2013–2017 forecast period, public IT cloud services will have a compound annual growth rate (CAGR) of 23.5%, five times that of the IT industry as a whole. Over the next several years, the primary driver for cloud adoption will shift from economics to innovation as leading-edge companies invest in cloud services as the foundation for new competitive offerings. The emergence of cloud as the core for new 'business as a service' offerings will accelerate cloud adoption and dramatically raise the cloud model's strategic value beyond CIOs to CXOs of all types. In Europe, we will already see a shift in 2014 with companies moving their focus from putting basics in place to using the cloud innovatively to differentiate businesses. Front-office will drive 50% of new apps investments in 2014 and new cloud applications will tend to integrate also mobility, big data and social technologies.

The cloud model has a dramatic influence on the tasks performed by IT departments as well as by IT vendors and service companies to produce and deliver IT changes. When companies buy IT as a service, whether from public cloud or a hosted cloud, they no longer need the technical skills and resources to build and maintain their own infrastructures and applications. Instead, they need to understand the available cloud services and technologies in the market, make sure users have access to the right services and applications, manage the usage and cost, manage data security and help users. In short, internal roles turn from IT builders and managers to brokers, strategists and business supporters. Over time a decline in demand for more traditional technology and operational support resources is expected (see Figure 1 below). Cloud centralises a very large part of the development effort at the provider, in theory freeing end-user organisations IT staff to do more value-adding activities rather than maintaining/fixing their on-premises applications and systems.

Compared to on premise vendors they compete against, cloud services vendors require more resources to develop, maintain and update their services. Cloud services have to be highly secure (they are a prime target for hackers), highly reliable (the minimum is 99.9% of the time, but providers are competing around this) and

sufficiently flexible to enable their customer to configure the services to their particular needs. However, consumers of cloud services have less need for software development. They do have to configure and administer the services they use with the tools and APIs provided by their cloud vendors. In the case of SaaS applications, this is very often done by non-technical "super user" administrators rather than IT staff. Some cloud customers also add their own custom code (in the case of SaaS, using the tools the SaaS vendor supplies so that the code can run on the SaaS platform). But customers do not have to (in fact they cannot) build, maintain/patch, update, customise, etc., the services themselves or the underlying systems that the service they are using runs on – those activities are entirely managed by the vendor.

Figure 1 Main trends of change of e-skills demand – Cloud Computing



Source: IDC 2013

New challenges on the horizon - the combination Cloud-Internet of Things the next frontier of disruptive innovation

A convergence of rapidly maturing technologies and markets is driving the development of the Internet of Things (IoT), which is finally moving from the "visionary" phase to real applications and services. IDC predicts

that in 2014 the number of installed intelligent communicating devices on the network will outnumber "traditional computing" devices by almost 2 to 1 globally. By 2020 the IoT is expected to reach a new level of maturity, based on fully deployed Machine-to-Machine (M2M) communication enhanced by cloud computing services and advanced networks. The combination of IoT and the dynamic compute capability of cloud has a huge potential impact on the deployment of new applications and services. However, there are relevant challenges to be met, particularly in Europe, to enable the evolution and successful diffusion of the emerging cloud-IoT ecosystem, ranging from creating favourable conditions for the development of new applications and services, to the development of standards and supportive regulation, to the availability of research and start-ups funding. To understand them one needs to look more closely into the interaction between IoT and clouds.

IoT and cloud are in many respects opposites – for example:

» IoT is based on decentralisation while cloud is based on centralisation
» IoT depends on device heterogeneity while cloud depends on the homogeneity of the environment
» IoT is mostly adopting event-driven architectures while cloud computing is essentially a service- oriented architecture.

However these differences make them complementary to each other. Cloud computing is expected to contribute to three core areas of the emerging IoT ecosystem:

» Applications – cloud is a natural place to build cloud-native analytical applications to create actionable insights (and actions) from data gathered from the IoT.
» Platforms and systems – cloud will provide a highly scalable (down as well as up, in short timescales) platform for applications and device management systems that are either conventional software applications or cloud native applications. It will also provide storage to capture potentially huge volumes of data generated by the IoT.
» Business analytics – cloud can provide flexible systems to analyse the data generated by the IoT.

Some the combinations of IoT and cloud are easy to anticipate, but many others will be disruptive and create new markets and business opportunities. Some of the most obvious combinations will involve using cloud to capture and store and data from an IoT, analyse the data, and feed the analysis back to inform and/or optimise processes back that the devices in the IoT are managing or monitoring. The reliability, capacity and ubiquity of wireless networks, coupled with increasing pervasiveness of wi-fi, enables real-time (or near real time) communication between IoT devices and the cloud. A growing number of examples already exist, such as crowd-sourcing of maps of radiation monitoring in Japan, a service launched by a group of device hackers on the commercial platform Xively after the Fukushima disaster, or in Smart cities and in the optimisation of self-driving vehicles (the connected cars). IDC will explore these potential trends and challenges on behalf of the EC DG CNECT in a new study focused on the development of a research and innovation policy leveraging the combination of cloud computing and IoT, whose results will be published at the end of 2014.

[1] http://www.idc.com/.

The Business Value of the Cloud and IT

Joe Weinman

Cloud computing, and related technologies such as big data, mobility, social, and the Internet of Things, represent a sea change in information technology, which is one of the most important technologies transforming the spheres of business, society, government, and the economy. Yet all too often, the focus on the cloud is technological and qualitative, rather than business-oriented and quantitative. Without a clear delineation of business benefits, technologies are ultimately doomed to irrelevance. Without quantification, all that remains is vendor hand-waving and supposition.

Who stands to benefit and how

Correctly characterising the business value and benefits of cloud computing and IT is an imperative for multiple sectors. For businesses considering adoption or growth of cloud initiatives, quantifying benefits leads to appropriate allocation of capital and executive attention. Moreover, IT's contribution to competitiveness via corporate strategy alignment is important in firm competitiveness, regional competitiveness and growth through job creation, infrastructure development, ecosystem clusters, etc., and thus guiding economic and regulatory policies.

Business models and sustainable services

Every vertical has different requirements and IT intensity, and within verticals, firms can be differentiated along a variety of axes. However, a useful model for consideration of cloud- and IT-based strategies is an extension of the Value Disciplines model formulated by Michael Treacy and Fred Wiersema [1]. They argue that companies can strategically differentiate based on Operational Excellence, Product Leadership, Customer Intimacy, to which we can add Accelerated Innovation. There are many variations of these strategies, which can be elucidated through an approach such as the Business Model Canvas [2]. These strategies can also often be applied to national or regional objectives. Cloud computing specifically and IT generally can invigorate these approaches [3]

Operational Excellence can benefit from the use of big data, mobility, and computing. For example, consider the use of data for process improvement, smart grid solutions to optimize electricity utilisation, or solutions to optimise physical logistics while maximising customer satisfaction.

Product Leadership can be enabled by the cloud and IT as well. Today's highly digital products can be digitalised and informationalised, and tied back to cloud-based services. Obvious examples include tablets and smartphones, which link back to cloud services such as app stores or functions such as search and social networking. Less obvious ones include connected cars, which link back to concierge, navigation, and entertainment services, wearable computing and biosensors, which may connect to healthcare or personal transformation services.

Customer Intimacy is shifting from an organisational approach to an IT-mediated approach. In many

sectors, rather than personal relationships and dedicated account teams, mass personalisation is enabling relationships at scale. Moreover, retailers, social networks, and entertainment firms are now using big data and sophisticated algorithms to increase revenues via upsell, cross-sell, and reduced churn, while also maximising customer value through effective recommendation engines.

Accelerated Innovation is being supported by contests, crowdsourcing, open source, and idea markets; all cloud-mediated. While the basic concepts date back exactly 300 years to the "Longitude Prize", today such markets and contests enable virtually anyone anywhere to contribute insights and creativity to solve problems. Even in the absence of a formal prize, access to information can open the floodgates of creativity broadly: consider the case of Jack Andraka, a high school teenager who may have developed a test for pancreatic cancer thanks to research he conducted on Google and Wikipedia. His technique may reduce the cost of testing by a factor of 26,000.

Measuring efficiencies and impact

In addition to qualitative strategies, many dimensions of cloud benefits can be quantified. For example, enhanced utilisation due to statistical multiplexing of independent workloads follows a reduction in penalty cost that is an inverse square root function [4], suggesting that even private clouds can have substantial benefits, and that mid-sized service providers or facilities can be viable. Under judicious assumptions, public clouds can offer economic advantages, but hybrid clouds generally show the greatest economic benefit, based on an economic breakeven analysis tied to variability statistics of the demand function, such as the Peak-to-Average ratio of workload demand [5].

As we move towards the Intercloud and the viability of cloud markets, the value of a cloud market in the presence of dynamic pricing and the absence of collusion can be quantified using the theory of order statistics as a harmonic sequence [6], again showing that even small or mid-sized regional markets can provide most of the benefit of an infinitely large one. The requirements for optimisation in large-scale distributed architectures also have important policy implications on wide area network regulation: anycast services can be best optimised by network services exposing information on congestion, and accepting controls for routing of specific flows [7]. These models can be made extremely rigorous, through the use of an axiomatic formulation of distributed, pay-per-use computing—an Axiomatic Cloud Theory [8] based on metric spaces, function spaces, set theory, measure theory, and Ð-algebras—and even used to prove computational complexity results, such as that cloud computing demand satisfiability is strongly NP-complete [9].

Links and References

[1] Michael Treacy, Fred Wiersema, *The Discipline of Market Leaders* (Addison-Wesley, 1995).

[2] Alexander Osterwalder,Yves Pigneur, *Business Model Generation* (John Wiley & Sons, 2010).

[3] Joe Weinman, *Cloudonomics: The Business Value of Cloud Computing* (John Wiley & Sons, 2012).

[4] Joe Weinman, "Smooth Operator: The Value of Demand Aggregation," http://joeweinman.com/Resources/Joe_Weinman_Smooth_Operator_Demand_Aggregation.pdf

[5] Joe Weinman, "Mathematical Proof of the Inevitability of Cloud Computing," http://joeweinman.com/Resources/Joe_Weinman_Inevitability_Of_Cloud.pdf

[6] Joe Weinman, "Quantifying the Value of a Cloud Computing Market," http://www.telx.com/blog/quantifying-the-value-of-a-cloud-computing-market/

[7] Joe Weinman, "Better Together: Quantifying the Benefits of the Smart Network," http://joeweinman.com/Resources/SmartNetwork.pdf

[8] Joe Weinman, "Axiomatic Cloud Theory," http://joeweinman.com/Resources/Joe_Weinman_Axiomatic_Cloud_Theory.pdf

[9] Joe Weinman, "Cloud Computing is NP-Complete," http://joeweinman.com/Resources/Joe_Weinman_Cloud_Computing_Is_NP-Complete.pdf



Challenges and future directions

Cloud-friendly laws?

Kuan Hon, Centre for Commercial Law Studies, Queen Mary University of London

Of EU laws affecting cloud computing, data protection laws are probably the most high-profile. Moves are afoot to modernise the Data Protection Directive, which did not envisage the massive scale of Internet use, let alone cloud computing. But how cloud-friendly are the current or proposed laws?

Who stands to benefit and how

All stakeholders, from cloud users, data subjects and regulators to cloud providers and other intermediaries, stand to benefit if data protection laws are updated in such a way as to uphold the underlying objectives while remaining technology-neutral.

Key concerns impeding the mainstream adoption of the cloud and new challenges on the horizon

Lawmakers, regulators and courts should concentrate on fundamental goals, the ends not the means, rather than seek to regulate technologies as such. Not all cloud services are the same, but laws which seem largely aimed at social media services risk threatening IaaS, PaaS and passive SaaS storage services alike.

A fundamental goal of data protection laws, the 'what', is protection of personal data – to which end EU laws regulate the **use** and **disclosure** of personal data. However, the 'how' has become confused. The 'how' is often considered to be access to personal data, including physical access. But physical access alone is neither necessary nor sufficient for access to intelligible personal data. This is because 21st century technology offers, and society can and often does employ, remote access to data, proprietary formats, distributed storage and encryption. Therefore, laws should focus on the true 'how': **control over logical access to intelligible personal data**. Such control can be technical (e.g. technical ability to access intelligible data remotely), and/or it can be legal (e.g. contract terms that ban or restrict someone with technical access from *actually* accessing data). This means two things.

Firstly, laws should regulate only those who have such control, and not those who don't. But many regulators don't take this approach in practice, leading to problems when trying to apply data protection laws to cloud computing. Thus, cloud providers (and sub-providers like IaaS/PaaS providers, maybe even datacentres) are considered to be 'processors' even if they hold only *encrypted* data, and have no idea whether personal data are included. Web hosting providers have the technical ability to access unencrypted data hosted on their infrastructure, but they have 'notice and take down' defences under the E-Commerce Directive, e.g. in relation to copyright infringement actions. Why shouldn't cloud providers, who may have technical access to unencrypted personal data that have been processed using their infrastructure, have similar defences to claims regarding personal data, based on knowledge of the nature of the data and control over the data?

Secondly, lawmakers and regulators should be concerned, not about physical access to data and geographical

data location, but about which countries can claim **effective legal jurisdiction** (irrespective of physical access or location) over **whoever can control logical access to intelligible personal data**. However, on many regulators' interpretation of the Directive, the **physical location** of personal data must be confined to the European Economic Area, regardless of encryption, with limited exceptions such transfers under the US Safe Harbour scheme. This approach restricts international data transfers unnecessarily, without focusing on the real issue: which countries have effective legal jurisdiction over those with logical access to intelligible personal data, regardless of the data's geographical location?

Another problem is that data protection laws regulate the use of processors based on 1970s outsourcing models, whereby data controllers hired processors to process data actively for them according to the controller's instructions, e.g. payroll processing. Under these laws, the controller-processor contract must oblige the processor to follow 'instructions'. But in cloud computing, providers don't actively process data as instructed by controllers; controllers retain direct technical access to their data, and process data *themselves* in self-service fashion using the provider's infrastructure. The underlying goal of rules regulating use of processors was to prevent **unauthorised use or disclosure** of personal data by processors or sub-processors. *That* is what laws should address, not 'instructions'.

It's also unclear whether EU data protection laws may apply to non-European cloud users if they use EEA cloud providers, EEA sub-providers (i.e. IaaS/PaaS providers) or EEA data centres to process personal data, or if they set up or acquire EEA subsidiaries – even where the data processed do not relate to EEA residents. These legal uncertainties may deter non-European cloud users from using EEA cloud providers or data centres even for backup purposes, and needs to be addressed.

The draft General Data Protection Regulation would make processors directly liable regardless of knowledge of the data's nature, including cloud providers. It would also let data subjects recover their entire loss from the cloud provider 'involved' if a controller, who happens to use cloud infrastructure, breaches data protection laws, e.g. by sending personal data to unauthorised parties from that infrastructure, leaving it to the provider to try to recover from the controller. It would not clarify the uncertainties regarding use of EEA cloud providers or infrastructure. It would restrict exports of personal data even further, and would no longer recognise the role that technological measures such as encryption can play in protecting personal data, instead recognising only 'legally binding instruments'.

This seems retrograde. The opportunity should be taken instead to focus data protection laws on the ends, not the means, so that we can end up with laws that are more future-proof, and not just cloud-friendly but tech-friendly.

Links and References

Centre for Commercial Law Studies, Queen Mary University of London, http://www.ccls.qmul.ac.uk/

Most of the above issues, and more including a case study of the UK public sector G-Cloud programme and its contract terms, are discussed in depth in chapters 3 to 5 of Cloud Computing Law (ed. Christopher Millard), OUP 2013, http://ukcatalogue.oup.com/product/9780199671687.do; Kindle edition, http://www.amazon.co.uk/Cloud-Computing-Law-Christopher-Millard-ebook/dp/B00GLO2OGW.

See also http://www.kuan0.com/publications.html, for Kuan's other publications and, for explaining cloud computing to the uninitiated:

12 C(haracteristic)s of Cloud Computing: a Culinary Confection, http://www.scl.org/site.aspx?i=ed26082 and 9 D(ifference)s of Cloud Computing, http://blog.kuan0.com/.

Further recommended reading

Chris Reed, Making Laws in Cyberspace (OUP 2012)
http://ukcatalogue.oup.com/product/9780199657612.do.

e-Infrastructure commons marketplace

Maryline Lengert - ESA, Bob Jones - CERN, David Foster - CERN, Steven Newhouse - EMBL-EBI

Researchers across Europe are looking for cost effective and sustainable IT services that can be combined to accelerate their work and increase its impact. Europe has a wealth of public and private sector service providers and when brought together they can create a ground-breaking open platform for innovation.

Who stands to benefit and how

Today, the majority of existing public e-infrastructures are supported by national/regional funding agencies and provide services that are free at the point-of-use. The financial support provided by the funding agencies is normally based on a fee linked to the cost of setting-up and operating a service rather than its level of usage. By introducing a pay-per-usage scheme as part of the overall funding model for the allocation of a fraction of the resources, as has been demonstrated within Helix Nebula, the funding agencies will have the information to be able to measure the level of usage of a service and whether it justifies their investments. In addition, implementing the pay-per-usage model will give some of the financial control to the users and they will favour those services which offer better value-propositions. The result of these changes to the e-infrastructure business model will reduce the total cost of service provisioning (processes building on digital data) and consequently contribute to their sustainability. The move to a federated marketplace model was described within the 'Strategic Plan for a Scientific Cloud Computing Infrastructure for Europe' [1] in what became the Helix Nebula Initiative and was generalised in 'e-infrastructure for the 21st Century' [2] issued by the EIROforum IT Working Group.

Introduction

The last decade has seen a tremendous growth in e-infrastructure and related activity in a number of research communities as a result of funding by the European Commission in the 7th (and earlier) Framework Programmes and by corresponding national investments. Consequent to this investment in capacity, research communities are presented with several individually excellent, but independent – cross-layer initiatives which present researchers with sometimes inconsistent technical approaches and disjointed managerial structures to achieving a production quality infrastructure. It is being widely recognised that this fragmented landscape has increased the complexity and reduced the willingness of research communities in their adoption of these e-Infrastructure services. Recently, a vision has emerged that addresses this fragmentation by proposing an 'e-Infrastructure Commons' [3], an open environment where researchers can flexibly discover and choose the services and service providers from either the public and private sector that they feel will best meet their needs.

Until recently, accessing research services has been a relatively closed static system with researchers applying to single local, national or European compute and storage service providers, usually through review process to receive (if successful) an allocation of resources on the designated systems. The advent

of publicly available commercial cloud services has provided an alternative approach for researchers and research communities. This approach has been further developed within the Helix Nebula initiative through the initial engagement of European Intergovernmental Research Organisations (EIROs), seeing it as a tool to perform generic data transformation processes. The Helix Nebula Science Cloud also brings unique data/knowledge/tools in a cross-domain market place catalysing science data to be seen in a different (unknown) context. Today science communities (earth, life, physics, etc.) want access and integration of many data sets regardless of location in order to address societal grand challenges.

Funding, sustainability and business opportunities

Today, the majority of existing public e-infrastructures are supported by national/regional funding agencies and provide services that are free at the point-of-use. The financial support provided by the funding agencies is normally based on a fee linked to the cost of setting-up and operating a service rather than its level of usage. By introducing a pay-per-usage scheme as part of the overall funding model for the allocation of a fraction of the resources, as has been demonstrated within Helix Nebula, the funding agencies will have the information to be able to measure the level of usage of a service and whether it justifies their investments. In addition, implementing the pay-per-usage model will give some of the financial control to the users and they will favour those services which offer better value-propositions. The result of these changes to the e-infrastructure business model will reduce the total cost of service provisioning (processes building on digital data) and consequently contribute to their sustainability. The move to a federated marketplace model was described within the 'Strategic Plan for a Scientific Cloud Computing Infrastructure for Europe' [4] in what became the Helix Nebula Initiative and was generalised in 'e-infrastructure for the 21st Century' [5] issued by the EIROforum IT Working Group.

Vision for an e-Infrastructure Commons Marketplace

The e-commons infrastructure marketplace, driven by the European Research Area, will provide public and private researchers with access to worldwide and world-class resources and services through a dynamic and sustainable marketplace. This overarching infrastructure, built on public and commercial assets, will cover the entire scientific workflow from research to production, from problem-solving to discovery and innovation. The marketplace will offer the broadest range of services available today and will participate in the development of those needed for tomorrow. It will ensure use of open standard and interoperability of service providers while adhering to European policies, norms and requirements.

To achieve this vision requires:

- More coherence and integration from services providers (public and private) in the e-Infrastructure Commons marketplace
- To engage researchers in all disciplines from all sizes of community
- To keep resources free at the point of use for researchers
- To link resource use to service provider income for sustainability
- To reduce the barriers to entry and simplify use for end-users.
- A holistic view of pan-European existing and planned e-infrastructure.

The marketplace should encompass both publicly funded and commercial assets so that the sum of these e-infrastructures, with all their complementarity and variety of "circles of influences", will create a new momentum in Europe, driven by science, to implement a knowledge-based society and economy.

Expected impact

The expected impact of this Marketplace is:

» Researchers, supported by large scale long term research infrastructure, drive the evolution of services for their research needs
» Funding agencies benefit from market forces to establish volume and price
» Create a fertile environment that nurtures new scientific ideas and challenges
» Service providers are able to attract revenues to sustain services
» It establishes an ecosystem that benefits downstream industry
» It assembles an ever growing marketplace building on Information as a service based on federation of data and IP meeting European security and integrity requirements
» It provides visibility and incentives to industry to invest in new assets (as a business case but also to use the science communities for testing cutting-edge technology as has been demonstrated by the CERN openlab project [6])

A governance and operational model will integrate and unify these services and stimulate expansion and adoption to new research communities, new service providers and the integration of new innovative technologies. The governance model shall involve all the stakeholders, including service suppliers and service consumers (end-users), as well as funding bodies seeking to use this platform as a policy implementation tool, to ensure that the market remains open and competitive.

Links and References

[1] http://cds.cern.ch/record/1374172/files/CERN-OPEN-2011-036.pdf.
[2] http://dx.doi.org/10.5281/zenodo.7592.
[3] http://www.e-irg.eu/images/stories/dissemination/white-paper_2013.pdf.
[4] http://cds.cern.ch/record/1374172/files/CERN-OPEN-2011-036.pdf.
[5] http://dx.doi.org/10.5281/zenodo.7592.
[6] http://openlab.web.cern.ch/becoming-sponsor.

The Future of Cloud in Europe – Successful adoption requires trusted clouds and BIG data

Tua Huomo, Future Cloud Action Line at EIT ICT Lab's

Cloud is a game changer for the European economy. Many services and applications have already become cloud-based and businesses and key infrastructures are becoming increasingly dependent on it. The momentum behind new innovations and business renewal include different types of services, applications and devices, which support the development of new digital service ecosystems. Overall, the speed of change in cloud technologies and services continues to be impressive and the main transformation is being led by the Cloud and Big Data integration.

Cloud-based business specifically including Internet of Everything (IoE) and Big Data Analytics are the top drivers of the European economy and society. Although Europe does not have a specific competitive advantage in cloud technology, due to its importance business and society, Europe cannot afford to rely on technology and services bought from elsewhere. The future of Cloud in Europe is promising but it requires secure multi-platform Clouds that users trust.

Europe's companies will benefit from a robust and efficient mission-critical Europe-based Cloud computing infrastructure that is established by leveraging Europe's strong position as trusted Cloud service provider. There are gains to be achieved in European competitiveness in the area of Cloud services and Big Data but European companies must be able to offer trusted Cloud services and solutions providing sufficient level of security and privacy to support business and personal requirements. Made in Europe solutions such as the unique Stratosphere platform, which is the only open source platform for Big Data analytics are examples of innovative new services being deployed by European companies for global markets. These are critical showcases of European Cloud innovation which should be recognised, trusted and applied globally.

Since 2009, EIT ICT Labs has set out to radically accelerate ICT innovation in Europe. A critical component of its 2014 – 2016 Strategic Innovation Agenda includes its Future Cloud Action Line which is focused on two critical priorities that will enhance mainstream Cloud adoption and European leadership on the global stage.

Priority 1: Establish leadership with trusted Multi-Cloud Infrastructures and Services:

EIT ICT Labs aims to leverage the strong position of the Europe as trusted Cloud service provider by showcasing Cloud services, infrastructures and solution offerings, best practices and examples, which address security, privacy and trust requirements across domains. Multi-cloud platforms will provide the basis for new innovative real-time Cloud services with the guarantee of the high level of security and privacy.

Priority 2: Cloud with Big Data Analytics:

EIT ICT Labs will facilitate the establishment of cross-technology communities through a European value-driven ecosystem and user community for Big Data in the Cloud. Business and application driven applied research and technology development will require cross-technology contributions (Cloud, IoE and Big Data). European solutions such as the Stratosphere platform are robust examples to strengthen Europe's position in the growing global Cloud service and solutions markets.



IEEE Intercloud Project – P2302 Working Group, and Global Testbed

David Bernstein & Joseph Weinman, IEEE Intercloud Project

Cloud Computing is a capability heavily utilised for research, Internet sites, and mobile telephony. However, unlike those utilities, clouds cannot yet federate and interoperate. This article reviews how the global Telephone System and the Internet were made interoperable amongst various Service Providers. The Intercloud Project [1] is described, where researchers tried several approaches to add interoperability to cloud computing. A technique which mirrors the way the Telephone System and the Internet was developed, which shows great promise, assuming that Cloud operators participate in implementing Intercloud protocols. It is shown that next generation services such as multiparty global video calling requires such a capability in order to deliver the best video quality possible. This article describes a "work in progress" – that is the process of the standards effort and test-bed efforts formed by the IEEE – and on-going activities. Finally, the technical highlights of the Intercloud architecture are described.

Who stands to benefit and how

Cloud computing is a new design pattern for large, distributed datacentres. Cloud computing offers end consumers a "pay as you go" model - a powerful shift for computing, towards a utility model like the telephone system or more recently the Internet.

However, unlike those utilities, clouds cannot yet federate and interoperate. In the Telephone Network, any phone can call any other phone with "direct dial". There is no requirement that the two phone users are connected to the same phone company! The phone network has even evolved, with Mobile, to allow a user to carry their phone to any country, "roam" with a provider, and make calls. This is amazing cooperation amongst telecommunications providers!

In the internet, any Internet connect browser can access any web site. The Internet Service Provider ("ISP") giving connectivity to the browser, does not have be the same Internet Service Provider hosting the web site. In fact, browsers can easily change ISPs – even in different countries – and web sites (as long as their name remains the same) can host in any location – and the system still works.

So far the global world of cloud computing does not have any of the capabilities of interoperability which have made the telephone network and the Internet such indispensable utilities.

As it turns out, early networks are never born with instant interoperability and federation. For example, telephone systems in different geographical areas did not interoperate at all, pre-arranged human intervention was required to manually plug together the phone systems of adjacent countries. International "direct dial" was not available until the relatively recently, in the 1970s.

As to the "on-line" word, in a precisely analogous evolution, the original online services such as AOL, Prodigy, and Compuserve had no interoperability between them. Content posted on one service could not be consumed by a client connected to a different service. Email could not be sent from a user of one service

to a user on another.

Interoperable solutions based on open standards implementation

The Multi-Cloud Approach

The first idea any research team comes up with for solving these kinds of problems is inevitably a "Multi-Cloud" approach. At first, one tries to solve the problem without changing any of the underlying clouds. Connections between clouds are made over the top via user APIs. In other words, the user places a mechanism - a box or a software API - in front of the multiple clouds (unbeknownst to the clouds themselves) which enables that user to view and use them all at once.

Let's look at our example of the Telephone Network. Does this Multi-Cloud approach enable direct dialling? How would one try to make "transparent" a number of phone companies? The common practice is to set up one box or service which is a member of each and every target phone network. When you want to call some number, you really call the box, tell it what number you want to get to, it decides what phone company is hosting that phone subscription, and posing as a "user" on that network, it connects you.

So while it looks like you have access to several phone companies, you don't have direct dial or "roaming". The modern day example of this is the Calling Card. In this way, you are using the "user APIs" of the phone system (phone numbers) to construct an over the top end to end connection.

We discarded the Multi-Cloud approach as insufficient; it just cannot provide the transparent interoperability which is needed. Because phone companies and ISPs decided to work together, they were able to do much better than Multi-Cloud for interoperability. They chose to change the networks to proactively federate.

The Federation Approach

In looking at how the interoperability problem was solved then, in the phone system and in the Internet, a theme emerged.

In each of these cases, special networking protocols were invented solve these problems. For the PSTN (Public Switched Telephone Network) a collection of protocols called the Intelligent Network ("IN") powered a new, out of band Signalling System (in its latest version called SS7), which allowed for transparent interoperability and federation, and paved the way for new features such as toll free calling, conference calls, call waiting, and network based voicemail, amongst other things.

For the Internet, a collection of conventions and protocols such as Autonomous System ("AS") numbering, Domain Name Service ("DNS"), Border Gateway Protocol ("BGP"), Simple Mail Transfer Protocol ("SMTP"), and Hypertext Transfer Protocol ("HTTP") (to name just a few) laid the groundwork for the global Internet on which many additional capabilities have been built.

Addressing new challenges on the horizon

We are quite confident that the world of Cloud will emerge just as the world of on-line services did, that is, based on the premise that interoperability is inevitable. Back in the on-line services days, users demanded interoperability because:

» They wanted to send mail from one on-line service to a subscriber on another. Interoperable email was

a huge driving force.

» Content owners did not want to put content in a proprietary form for each on-line service. The idea of a "web site" that anyone could access was a revolutionary idea.
» Users did not want to have a separate client for each service. The idea of a "universal browser" was another powerful idea.
» Users wanted to be able to search "everywhere". Remember the AOL Keyword? This was useless on Compuserve or Prodigy. An interoperable network gave birth to internet-wide wide search.

Back then, interoperability begot the "Internet"; in the future, interoperability will cause the emergence of the "Intercloud".

By way of analogy, the airline industry, broadly speaking, comprises more than just airlines; it includes travel agents, air travel web sites, charter operators, travel insurers, food and fuel service companies, tour operators that bundle and/or white label carriers, airport concessions, and so forth. Perhaps more subtly, even if none of those third parties existed, airlines would still want or need to be able to codeshare, coordinate air traffic to minimise fuel costs and collisions, and efficiently transfer passengers and cargo to each other. The cloud computing industry is not much different.

Although now perhaps taken for granted, multilateral organisations such as the International Air Transport Association and International Civil Aviation Organization were and are needed to propose and gain consensus around mechanisms such as standard addressing for resources (airport codes such as SFO, LAX, and JFK), service provider identification (airline codes such as UA and BA), and tower communications (the English language and the NATO Phonetic Alphabet, i.e., Alpha, Bravo, Charlie, etc.). Some airlines, such as those in the STAR alliance, are federated, offering a tighter level of integration, e.g., cross-airline frequent flyer perks. Similarly, cloud service providers, customers, and third parties are likely to benefit from a variety of standards, interoperability, communications, and federation mechanisms.

One cloud provider might be able to virtually extend its physical footprint (think airline code sharing) by leveraging investments in facilities made by another provider, even dynamically provisioning computing resources in real time.

A cloud provider suffering from an outage or insufficient capacity might be able to transition existing customers to a competitor, the way one airline might rebook a passenger with a competitor if experiencing a maintenance issue, flight cancellation, or overbooking.

A customer would easily be able to run an application requiring more than one provider, in the same way that a passenger traveling from New York to Seoul might combine services from United Airlines and then Korean Air, transferring at a neutral interconnection point such as Tokyo's Narita airport. In the compute world, carrier neutral interconnection and co-location facilities are often the "airports" that enable such handoffs.

Customers would be able to compare cloud providers total offer and select one with the best combination of price, performance, availability, and so forth, the way a traveler might use Travelocity to trade off total price including ticket, taxes, and baggage fees, departure time, arrival time, number of stops, total travel time, and reliability (e.g., % on-time arrival), etc. Even a simple directory listing service providers that could meet requirements would be helpful; in the same way that it may be necessary to determine which carriers fly to, say, Istanbul.

Cross-provider bundles and workflows would benefit from Intercloud standards and mechanisms as well. Today, tour operators ("aggregators") put together hotel, airline, train, bus, and cruise ship services into a single bundle (say, "The Mediterranean Like You've Never Seen It Before") and manage the flow of customers across those services (the helpful people guiding travellers from baggage claim to the bus that takes them to the cruise ship). Complex information technology applications and "mash-ups," requiring, say, credit card validation and billing, scanned-image-to-text optical character recognition, secure messaging and the like, could be composed from base services offered by a variety of best-in-breed operators.

Even if not bundled, virtual operators could easily resell physical cloud provider services, perhaps offering a greater degree of intimacy or solution engineering than core providers might want to. This is not dissimilar from a human travel agent selling a ticket. After all, they don't actually run an airline, but can still offer value-added guidance.

Cloud providers will likely benefit from greater customer flexibility to migrate workloads and data to competitors, by eliminating fear of lock-in as a barrier to cloud adoption. Just as airlines (try to) ensure proper baggage handling, "reliable application transport" in the Intercloud will ensure that data and applications are correctly transferred across providers. In the same way that online ticketing sites and interline baggage handling enrich the usability of airline services, rather than supplanting them, the Intercloud offers the promise of enhancing the customer benefits of the cloud by complementing and facilitating cloud vendor products and provider services.

[1] http://cloudcomputing.ieee.org/intercloud. See also, http://www.intercloudtestbed.org/

Expert insights on data regulations, contracts and security

Contractual and data protection aspects of cloud computing

Paolo Balboni & Domenico Converso, ICT Legal Consulting

ICT Legal Consulting is a law firm with offices in Amsterdam (International Desk), Milan, Bologna and Rome and with legal experts operating in fourteen countries worldwide. The Firm offers unique expertise in Information and Communication Technology (ICT), Intellectual Property (IP), Privacy and Data Protection Law.

ICT Legal Consulting has been involved in several important European cloud-related projects and we assist numerous multinational companies in negotiating and drafting cloud computing agreements (both for customers and providers). Our main goal is to turn legal advice into strategic advice to create competitive market advantages.

Addressing key concerns

Cloud computing services have gained considerable momentum in recent years. However, many companies, especially small and medium-sized companies, continue to express concerns in terms of contractual issues and privacy management.

Whenever a cloud computing service is highly complex, it naturally implies the presence of a highly complex agreement, as well as the management of a large amount of data, often personal and sometimes sensitive. To date, the standardisation and use of adhesion contracts has been the norm by cloud providers. Such an approach tries not to leave space for negotiating on the part of their customers, often with little information about technical and organisational measures in place to guarantee security and confidentiality of data processing.

However, the trend seems to be changing as providers begin to realise that to be competitive they need to be more aware of customer concerns, bringing more flexibility into negotiations and more willingness to demonstrate the contractual and technical robustness of their services.

For these reasons, concerns such as the structure of a cloud computing agreement, confidentiality between parties, the presence of multiple parties and sub-contractors, the balance of responsibilities, the choice of applicable law or the security of data and compliance with privacy and data protection regulations, represent essential elements that need to be seriously taken into account both on the side of the cloud provider and cloud customers.

In order to address some of these concerns, ICT Legal Consulting hosts innovative workshops and takes part in panel debates with the aim of sharing experiences, strategic tips and legal recommendations.

Who stands to benefit and how

ICT Legal Consulting workshops and panels focus on sharing practical legal tips and stimulating debate on aspects related to contractual issues and data protection, an innovative approach that has already received particular interest in a number of EU countries.

These events offer innovative companies, European SMEs and IT professionals a practical and concrete legal overview on cloud computing services, looking at them from a contractual and privacy perspective.

Expert panellists focus on a twofold objective:

» Highlighting the main contractual clauses and privacy legal issues involving cloud-based services.

» Pointing out the main recommendations that need taking into account before entering into a cloud service agreement.



Service Level Agreements for data protection and data security

Wolfgang Ziegler, Fraunhofer SCAI

The Grid and Cloud middleware research group at the Fraunhofer institute SCAI is a key player in the definition of standards for electronic dynamic service level agreements (SLAs). Besides the specification of SLA negotiation and creation, we have a focus on the languages needed to express service description terms, service level objectives and key performance indicators (KPIs) within these SLAs. For example, in the European project OPTIMIS [1], we developed a solution for including standard contractual clauses (SCC), binding corporate rules (BCR) and protection of intellectual properties (IPR) in a service level agreement.

Who stands to benefit and how

Increased data protection and data security through binding service level agreements will be beneficial for all parties involved: the infrastructure providers, service providers and end-users. The benefit of the end-user is most evident as the end-users (a company or an individual) using cloud infrastructure for their businesses may request a service with a defined level of data protection and reach a binding agreement on this before actually using the service. This also applies to the general public using service offerings that are based on SLAs between service provider and infrastructure provider that include data protection and data security clauses.

Addressing key concerns impeding the mainstream adoption of the cloud & the need for standards

One of the big concerns around public cloud infrastructures is the achievable and realistic level of data protection and data security. Likewise, how can a level be reached and guaranteed that satisfies the real company or personal protection requirements? What's more, there are governmental or state requirements that companies or public organisations need to take into account when storing or processing data in a public cloud.

Today, most SLAs between a large service provider and a customer are clearly in favour of the provider when it comes to liability and compensation. Providers of public clouds only offer limited and non-negotiable SLAs should any issue arise. With regard to data protection requirements, at best customers can only select a region where their data should be stored and processed. However, there is no guarantee and, even worse, there are no mechanisms for the customer to monitor and control the actual geographical location of the resources provided. Thus, for most companies and pubic administrations, the use of public cloud resources is close to impossible.

To overcome this situation, we need both new forms of SLAs between providers and customers and means for the customer to verify the state of an SLA while using the resources.

Dynamic electronic SLAs between a customer and provider could change the current situation as by allowing the shift beyond the immutable and biased SLAs of the providers, with more flexibility to adapt the SLAs to the need of the customers. Clearly, the creation and negotiation of dynamic electronic SLAs must be based on standards to achieve interoperability and to empower the customer to compare the offerings of different cloud providers, and, where necessary, leverage targeted tools, interfaces or cloud brokers. In addition to this, standardised languages for expressing service description terms, service level objectives and KPIs are needed to request and negotiate SLAs covering the same service levels from different providers before choosing the best provider.

Naturally, more detailed SLAs also require the means to verify whether any aspect of the SLA risks being violated during infrastructure service usage. While this monitoring is usually done by the provider to take appropriate counter measures against risks, no monitoring information is available for the customer unless a customer sets up its own dedicated monitoring when using the resources. What is needed here is a monitoring interface offered by the provider, which is trustworthy for both the provider and the customer. Another approach is providing the interface to monitoring data through a trusted third party.

Finally, a mechanism is needed to enable the customer to verify the geographical location of the resources provided. This is probably the most difficult part as it requires combining multiple technologies like certification and different measurements. Consistently, solutions have yet to be developed.

[1] www.optimis-project.eu.



Cloud computing in the public sector – A European Perspective

Dimitra Liveri, ENISA

Neelie Kroes, Vice President of the European Union, has frequently highlighted the potential benefits of cloud computing: "*Cloud Computing will change our economy. It can bring significant productivity benefits to all, right through to the smallest companies, and also to individuals. It promises scalable, secure services for greater efficiency, greater flexibility, and lower cost*".



Cloud adoption trends

Public and private sector organisations are increasingly adopting cloud computing in a shift away from in-house IT approaches towards outsourcing to large cloud service providers. Evidence suggests that in a couple of years, around 80% of organisations will be dependent on cloud computing.

Public administration is playing a key role in the uptake of cloud computing to gain the benefits of scalability, elasticity, high performance, resilience and security, together with cost efficiency, which make it an attractive business model for public bodies. In the same light the use of cloud services could enable and simplify citizen interaction with government by reducing information processing time, lowering the cost of government services and enhancing data security. Governmental Clouds offer to the public bodies, including ministries, government agencies and public administrations (PAs), the potential to manage security and resilience in traditional ICT environments and strengthen their national cloud strategy.

Key concerns impeding mainstream adoption

However, the adoption of cloud computing also raises concerns about security and associated risks. The European Networking and Information Security Agency (ENISA [1]) has published a number of studies on these concerns, providing guidance on how to procure cloud services securely. In 2013, ENISA conducted a study, focusing on the implementation of cloud technologies in the public sector, covering 23 countries both in and outside the European Union [2]. This research shows that there are still concerns to be addressed before cloud adoption by public authorities can become mainstream. Chief among the barriers are security and privacy in the cloud, preventing public authorities from moving to the cloud. The same risks were also

identified in the 2009 ENISA risk assessment [3], offering further evidence that concerns about the risk of loss of control/governance and data locality remain high on the European adoption roadmap.

Priority areas

Efforts are underway to find ways of mitigating these risks. It is important that such efforts take place on a national or even pan European level. Priority areas include the development of national cloud strategies to foster the adoption of government clouds, the development of a common framework for Service Level Agreements (SLAs) focused on government clouds, the build-up of a certification framework for cloud providers, the adoption of measures to ensure security across both private and public deployment models. The European Cloud Computing Strategy (September 2012) focuses on addressing concerns by defining a core set of actions, such as voluntary certification schemes and identifying necessary standards.

Links and References

[1] ENISA, http://www.enisa.europa.eu/.

[2] Good Practice Guide for securely deploying Governmental Clouds ,
http://www.enisa.europa.eu/activities/Resilience-and-CIIP/cloud-computing/good-practice-guide-for-securely-deploying-governmental-clouds.

[3] Cloud Computing Risk Assessment, ENISA, 2009,
http://www.enisa.europa.eu/activities/risk-management/files/deliverables/cloud-computing-risk-assessment.

Negotiating cloud contracts - from both sides now

Kuan Hon, Centre for Commercial Law Studies, Queen Mary University of London

Key concerns impeding the mainstream adoption of the cloud and new challenges on the horizon

Cloud computing services are generally offered on cloud providers' standard terms. 'Off the shelf' cloud services are easy and quick to procure, just by clicking through and (in the case of paid services) providing credit card details. This has contributed to the growth of shadow IT use, where IT, procurement and other departments may not know the extent to which employees of the organisation are using cloud services, including for confidential and personal data.

The Cloud Legal Project [1] at the Centre for Commercial Law Studies, Queen Mary University of London [2] conducted ground-breaking research on cloud contracts, surveying some 30 sets of standard cloud contract terms in 2010, and analysing the negotiation of cloud contracts through anonymised interviews with cloud market players in 2012.

With providers' standard terms, key specific risks identified by the research related to wide liability exclusions and disclaimers, sub-contracting by providers (e.g. SaaS built on IaaS or PaaS), ability for the provider to change or discontinue the service at any time, and recovery of data after termination. It was questionable whether some of these terms were legally enforceable, particularly against consumers or under laws regarding unfair standard terms.

Users may seek changes to providers' standard terms for several reasons. Terms tend to favour providers, unsurprisingly – although not always, e.g. the terms of providers with a legacy of enterprise rather than individual consumer customers. There are commercial reasons, notably SLAs and risk allocation. And the user needs to remain compliant with laws and regulations affecting it when using cloud, in particular data protection laws regulating processing personal data, and financial services regulation.

Whether users can negotiate successfully depends as always on the user's bargaining position. Even huge corporates have had difficulty persuading large providers to agree to any changes. Our research found that users in the best position to secure changes tended to be financial institutions and government/public sector users. Most of these contracts are confidential but some have been published. The provider's position is also relevant – smaller providers, unsurprisingly, were more willing to negotiate. Integrators can play a significant role too, sitting in the middle and contracting with both user and provider. Some integrators have proved willing to accept liability desired by the user but rejected by the provider, but of course this leaves the integrator exposed to the risk mismatch.

In our negotiated contracts research, the top 6 points most negotiated (which sometimes proved to be 'deal breakers') were exclusion/limitation of liability, SLAs, security/privacy, lock-in and exit, providers' rights to modify the service unilaterally, and intellectual property rights. Liability was far ahead of the others. There are, of course, several key tensions here. 'Guaranteed' liability and security may be possible, but will cost

money, and seems at odds with the model of cheap or free public cloud. The biggest providers may end up the winners as they are the most likely to be able to control the entire supply chain, from datacentres to any IaaS/PaaS layers, and therefore be able to offer the guarantees sought particularly by users in regulated sectors.

There is still some way to go in improving user awareness and educating users about the risks of using cloud computing and ways to mitigate their risks, whether technical or legal. More guidance and pre-contract risk assessment checklists for users would assist, and users should be encouraged to implement backups and encryption whenever possible and appropriate. User demand may push, and indeed show signs of having pushed, providers to make their terms more customer-friendly for market competitiveness reasons.

As for laws and regulation, it needs to be borne in mind that imposing direct liabilities on all cloud providers alike, without regard to whether they can or do access and use or disclose intelligible data, may drive prices up and discourage infrastructure providers in particular from offering their services to EU customers. The market seems too varied to prescribe standard terms for cloud contracts; consumer protection laws do exist, and can and have been used. Certifications, codes of conduct and seals seem promising, but the devil will lie in the details of such schemes, and incentives will be needed to persuade providers to invest in obtaining certifications etc., such as liability reductions or defences for those who have complied with such certifications. We may see the emergence of a 3-tier cloud, with free or cheap cloud services, more expensive services certified as 'fit for personal data', and even more expensive, highly-secure cloud services e.g. for financial services.

Links and References

[1] http://cloudlegalproject.org.

[2] http://ccls.qmul.ac.uk.

Most of the above issues, and more including a case study of the UK public sector G-Cloud programme and its contract terms, are discussed in depth in chapters 3 to 5 of Cloud Computing Law (ed. Christopher Millard), OUP 2013, http://ukcatalogue.oup.com/product/9780199671687.do;

Kindle edition, http://www.amazon.co.uk/Cloud-Computing-Law-Christopher-Millard-ebook/dp/B00GLO2OGW.

See also http://www.kuan0.com/publications.html, for Kuan's other publications and, for explaining cloud computing to the uninitiated,

12 C(haracteristic)s of Cloud Computing: a Culinary Confection, http://www.scl.org/site.aspx?i=ed26082 and 9 D(ifference)s of Cloud Computing, http://blog.kuan0.com/.

Smart services, smart society – Perspectives from the public sector

EDM an Austrian eGovernment CLOUD Services that protects the environment and makes complex processes manageable – EuroCloud 2013 Best Cloud Service Use Case Public Sector

Franz Mochty, Federal Ministry for Agriculture, Forestry, Environment and Water Management, Austria

EDM is an award-winning eGovernment tool developed by the Austrian Federal Ministry for Agriculture, Forestry, Environment and Water Management in co-operation with the federal provinces over a number of many years [1]. EDM is extremely extensive and complex Software as a Service (SaaS) application, comprising a network of 22 applications dealing with various environmental requirements. EDM makes complex legal provisions manageable through menu guided processes and automated validation. The objective of EDM is to create clarity and legal certainty for all stakeholders by supporting a uniform application of Austrian and European legislation in the environmental sector, thus making a major contribution to maintaining the high standard of Austrian environmental protection.

Who stands to benefit and how

EDM covers the whole environmental sector. For instance, EDM applications handle the complete waste management processes. The many obligations of companies involved in the generation, collection or treatment of waste - such as those concerning the recycling of used electrical equipment under the polluter pays principle - are supported by EDM. The following organisations benefit from using EDM:

- Federal Ministry for Agriculture, Forestry, Environment and Water Management and other federal ministries.
- Authorities of all 9 federal provinces.
- District authorities across Austria.
- All companies involved in the generation, collection or treatment of waste.
- Industries and businesses producing emission to air and water.
- All citizens through better environmental protection.

Around 1,500 employees in various administrative authorities work with EDM. EDM is used in 4 federal ministries - Health, Finance, the Interior and the Federal Ministry of Agriculture, Forestry, Environment and Water Management - as well as in all Austrian provinces and all 95 district authorities.

Around 800,000 reports a year are submitted in EDM, consisting in part of several thousand datasets, with a total of over 20 million hits of EDM services annually.

EDM is a cloud application in the classic sense of the NIST definition. EDM can be accessed on the Internet via standard protocols, is not bound to any specific terminals nor does it require any local installation. All resources and data are freely available for the many users in the form of a pool. Of course, being an eGovernment application means that there is a legal obligation that the data is stored exclusively in

Austrian datacentres. The services of EDM are provided free of charge and without requiring a licence. Access and export of data are carried out in compliance with the strict requirements of Austrian data protection legislation. There are, however, a series of web services and XML interfaces available for import and export, so that highly automated processes can also be supported by EDM.
The clearly arranged, well-structured and in part menu-guided design of EDM makes even complex processes manageable. Without electronic support it would be almost impossible to manage the many requirements in practice. Environmental inspections, for example, can be carried out efficiently with less administrative burden for companies and authorities.

Better services, smarter society

EDM is fully integrated into the Austrian eGovernment environment, e.g. information of companies in the Austrian business register is used directly and is an integral part of the Austrian Portal Group. A basic EDM principle ensures that data is collected and managed only once, namely when they first arise and after that are transmitted and processed exclusively without media discontinuity.
EDM has a cross-administrative design. So EDM enables the integration of authorities at different administrative levels and with different areas of competence. This is important because there is often more than one authority responsible for a specific permit or report.
The EDM programme is used actively in several European and international work groups preparing the way for and promoting interoperability with other European and international eGovernment systems.
Co-operation between waste and economic and business authorities has been improved Austrian wide thanks to the new implementation of the Directive on Industrial Emissions. EDM includes environmentally relevant information concerning 45,000 registered companies with 17,000 locations and 20,000 plants.
Among other things EDM supports environmental protection. Environmental protection boosts the economy and creates new jobs. The introduction of EDM has made it more difficult for companies to disregard environmental provisions. This prevents environmental pollution and helps with imposing the polluter pays principle. The implementation of EDM as a cloud service gives flexibility for adaption to legal developments and the gradual expansion of user support.

Links and References

[1] http://www.eurocloud.org/congress/.

Cloud for Europe – Challenging the European market for public administrations

Linda Strick, Fraunhofer FOKUS

The Cloud for Europe project [1] brings together industry and the public sector to provide fair conditions for a digital single market for cloud computing in Europe. Cloud for Europe will give a clear view on public sector requirements and usage scenarios for cloud computing. The project addresses the objectives of the European Cloud Partnership and contributes to adopting a well-defined European Cloud Computing Strategy for the public sector. Joint pre-commercial procurement (PCP) will be used as an instrument for promoting, among European private sector players, innovative solutions for cloud services that best fit the public sector needs.

Who stands to benefit and how

Lessons learnt from the pre-commercial procurement process will be transformed into best practices and recommendations for future pre-commercial procurement activities. This includes establishing suitable contractual terms and conditions for future cloud procurements. Guidance and training materials will be produced for public authorities who would like to apply a pre-commercial procurement scheme to procure cloud services. Industries, especially SMEs, are invited to participate in discussions about the potential of future public sector cloud solutions. As a result of the pre-commercial procurement process, awarded industry will have a prototypical implementation of a pre-product that can be transferred into a product and offered to the all public administrations. The conditions are set by the pre-commercial procurement process.

Interoperable solutions based on open standards implementation

Interoperability is a basic requirement for cloud services that are shared between public authorities or across borders. It allows fair competition and is fundamental to integrating components, so that an ongoing innovation process is possible. But interoperability demands common technical and legal parameters, which are related to open standards and governance.

Addressing key concerns impeding the mainstream adoption of the cloud

Transparency enables trust and encourages active participation in political decision-making processes, supporting co-operation within public administrations and with industry. Furthermore, it fosters continuous innovation. Cloud computing facilitates transparency through shared platforms, standards and principles. From a political perspective it requires common minimum denominators in key regulatory areas to unleash cloud in the public sector and in sensitive areas in the private sector. Regulatory decisions at the political

level can help to overcome the obstacles. Common minimum denominators in key regulatory areas are required to unleash cloud in the public sector and in sensitive areas in the private sector.
Data protection and security are the obstacles hindering the public sector to go for cloud computing. Strong contracts are needed to ensure that any breaches of data security are avoided and detected. Contract templates, codes of conduct, and model clauses could be provided to procuring partners to manage and build trust in cloud computing. Measurements can be used as a basis to check the security in the cloud.

Links and References

[1] www.cloudforeurope.eu.

Funding

Cloud for Europe is funded under the European Commission's 7th European Framework, Software & Services, Cloud, DG CNECT (FP7 ICT Call 10).

Me and My Cloud

Interoperability is the key to freedom in the Cloud

Michel Drescher, EGI.eu

The European Grid Infrastructure (EGI) [1] is building a federated, standards-based IaaS Cloud platform, building on its decade-long experience in delivering a reliable, federated Grid infrastructure for scientific computing and e-Research across Europe and worldwide. Ultimately, the cloud solution offers advanced ICT capabilities for research, virtualised resources to run any environment chosen, cloud storage for easier sharing of data, and a number of support services to ensure applications run as efficiently as possible. Federation is enabled by a set of core services such as seamless authentication and authorisation of users, gathering of accounting information, information discovery, monitoring and VM management across multiple cloud domains. Federated cloud providers engage by establishing an SLA with the federating organisation – EGI.eu.

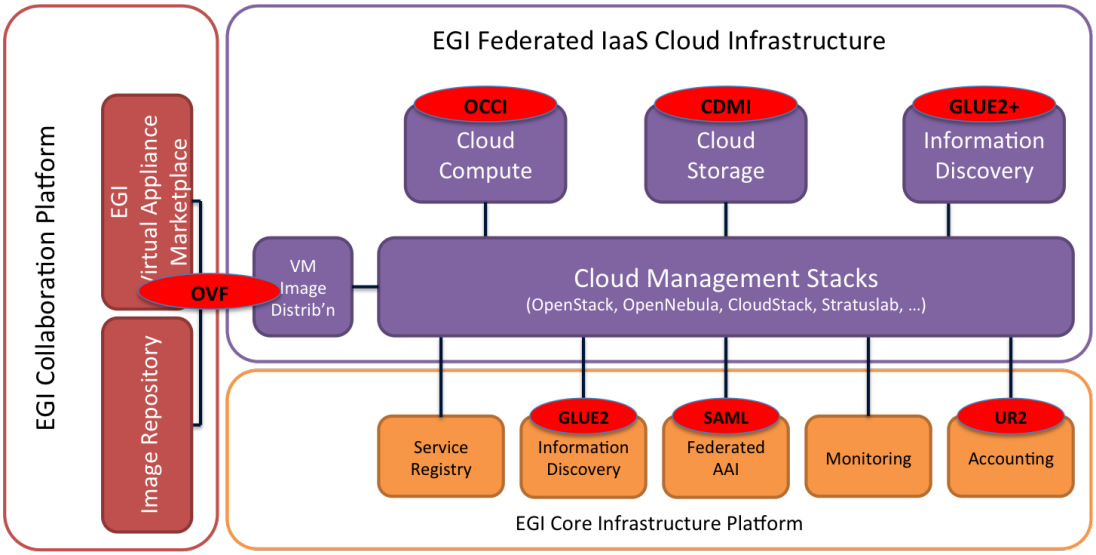
EGI chose a federation-based governance model, balancing out the individual freedom of participating suppliers, and the cost benefits of providing common services once instead of over and over again. Technical consistency in the service delivery between participating suppliers is ensured by extensive use and mandating of publicly defined interface specifications such as **OCCI** [2], **CDMI** [3] and **OVF** [4]. At the same time, customers enjoy a single point of contact for common services such as accounting, resource monitoring, etc. EGI's service catalogue [5] and solution portfolio [6] allows researchers to make use of EGI's services independently through one of the largest aggregation of federated IaaS Cloud resources in Europe. EGI's federated Cloud Infrastructure Platform will go into **production in May 2014** with an initial capacity of 2,000 cores and 15 TB of storage**.** By the end of 2014, EGI will provide to its customers more than **10,000 cores** and almost **1.5 PB of storage.** Beyond that, EGI's capacity building programme will include both increasing individual member capacity, and integrating more resource providers, in order to reach its goal of providing **10M cores** and **1 Exabyte Cloud storage** for its customers by 2025.

Who stands to benefit and how

Key to EGI's IaaS Cloud federation is its design as an *enabling solution* for its stakeholders, much like virtualisation being an enabling technology for IaaS Clouds. With this in mind, EGI is targeting large research communities (or partnerships of communities) supporting multidisciplinary science in Europe and beyond. EGI Distributed Competence Centre provides support to research communities to develop their platforms on the federated cloud, offering through its federation both academic and commercial cloud resources, seamlessly integrated, supporting a variety of Cloud infrastructure requirements at the same time.

Interoperable solutions based on open standards implementation

EGI strongly promotes **open standards** and, consequently, **interoperability**, as these enable a fair, open, transparent, and level playing field for all participants, both on the demand side and the supply side. EGI does not subscribe to the notion of standards stifling innovation and competition; instead, standards are seen as a mechanism to regulate governance and control over access to a market, where supply side players can compete and excel through their means of service delivery. Standards also allow the consumers of a service to compare competing offers, and ease risk management, such as reducing or eliminating the risk of a single point of failure when engaging with only one Cloud service provider – think of data centres buying connectivity from several competing network providers.



EGI federated Cloud Infrastructure Platform architecture and standards

EGI has chosen the following set of standards as being part of the mandate to participate in the Cloud infrastructure federation:

- **OCCI [2]:** A family of specifications defining access and management operations for IaaS Clouds. Highly extensible hence not limited to IaaS only. Extensions for business services such as accounting & monitoring are currently discussed.
- **CDMI [3]:** A specification defining a generic management interface, metadata management, and data access protocol negotiation for Cloud storage services. A near-perfect complement to OCCI (see above).
- **GLUE2 [9], GLUE2+:** Information model specification. Originally defined for academic Grid resources, an extension for Cloud resources is currently being discussed.
- **SAML [10]:** Specification regulating details of authentication and authorisation.
- **UR2 [11]:** Similarly to GLUE2 originally designed for Grid resources but includes in its latest published

revision definition of accounting records for Cloud resources.

» **OVF [12]:** A specification for a container structure that includes binary VM images and basic deployment and contextualisation instructions. Starting point for managing virtual appliances.

This approach is very much in line with the latest Future of Cloud Computing survey [7] but even more so with the Digital Agenda for Europe:

> *"As science is increasingly driven by the processing of big data, researchers need access to science clouds and other e-infrastructure that satisfy their requirements. Interoperability of services and applications is a key concern because it broadens choice and ensures a level playing field for both service providers and users, driving competition and innovation. I am a fervent defender of interoperability and it is a priority of the Digital Agenda for Europe."*
>
> **Neelie Kroes, Siena Roadmap, June 2012**

Business models and sustainable services

While based on standards, the EGI federated model allows the offering of services through any other IaaS access interfaces including proprietary ones. By aiming for a lightweight, yet comprehensive federation framework and composition, it allows its members to find their own sustainability means through compatible business models, which complement their own strategy by either strengthening existing service uptake, or broadening their service catalogue.

Different marketplace models are being evaluated. Services may be delivered through a marketplace tailored to the needs of the EGI users, or through the integration into a larger service, e.g. the Helix Nebula Marketplace [8].

Addressing new challenges on the horizon

Building capacity not only matches growing demand in cloud computing, but also ensures that new use cases and capability demand will be properly matched in due time. The EGI Federated Cloud will address long-standing issues in cloud computing, and support future challenges such as the Big Data problem: While each individual Cloud provider has a wider variety of solutions at hand to provide networking services to their customers, inter-provider networking is an entirely different problem. The EGI federation will tackle this by pushing the limits to provide lightpaths as a service between members of the federation. At the same time, the amount of data generated, stored and curated is growing beyond imagination – data develops a sense of gravity, which gives rise to the need of flexibly placing compute services in close proximity to (practically inert) data resources, automated, on-demand and accounted for.

Part of EGI's federation model is its extensive network of collaborations and pioneering researchers. EGI's service catalogue and solution portfolio allows researchers to make use of EGI's services independently through one of the largest aggregation of federated IaaS Cloud resources in Europe, or by focused strategic support and collaboration in projects through its community networks and support & community driven innovation solutions, charged through the federation's governance bodies.

Links and References

[1] https://www.egi.eu/.

[2] http://occi-wg.org/about/specification/.

[3] http://www.snia.org/cdmi.

[4] http://www.dmtf.org/standards/ovf.

[5] http://www.egi.eu/services/.

[6] http://www.egi.eu/solutions/.

[7] http://go.egi.eu/AnnualCloudSurvey2013.

[8] http://www.helix-nebula.eu

[9] http://www.ogf.org/documents/GFD.147.pdfÐ.

[10] https://www.oasis-open.org/committees/tc_home.php?wg_abbrev=security.

[11] https://www.ogf.org/documents/GFD.98.pdf.

[12] http://www.dmtf.org/standards/ovf.



The evolution of the ~okeanos IaaS cloud service

Vangelis Floros, GRNET

GRNET is the Greek national research and education network (NREN) responsible for the provision of advanced e-infrastructure services to its user base. To serve this community, GRNET has developed its own IaaS cloud solution, which is providing virtualised computing resources free of charge to Greek universities and public research centres that already utilise GRNET's network infrastructure and services. The cloud service, called ~okeanos, is built on top of existing proven open source software (e.g. Google Ganeti) which has been expanded in house in order to provide a robust and complete IaaS cloud solution. After a long period of testing, ~okeanos has recently moved to beta production phase and is already offering cloud resources to thousands of users.

Who stands to benefit and how

The service was initially conceived and designed with the Greek research and education community in mind, that is, the natural user base of GRNET. However, it soon it became evident though that there is a wider potential for usage in a broader environment. With this in mind, ~okeanos joined the EGI Federated Cloud activities with the aim of enhancing its interoperability features and enable the offering of computing resources to the high-productivity federated infrastructure offered by EGI. Moreover, ~okeanos has recently opened access to the GÉANT community through the ~okeanos GLOBAL initiative. The launch of ~okeanos GLOBAL has been enabled through GÉANT's GN3+ SA7 Support to cloud activities. By using the eduGAIN interfederation service, users across the globe can freely access and try ~okeanos. Last but not least, ~okeanos is one of the cloud services participating in the European CELAR project, which is developing advanced elasticity capabilities for IaaS clouds.

Innovative, interoperable and sustainable services

~okeanos is an excellent example how NRENs s can offer innovative, state-of-the-art public services through access to Structural Funds. The development of this kind of service has various positive implications in the GRNET business portfolio. First of all, brings added value to the academic network. With the support of a national funded project, this has attracted thousands of users, be they students, professors or researchers from various Greek institutions, who have used the service in the context of numerous professional, academic but also personal applications.

The increased demand for computing resources has pushed GRNET to expand its underlying computing capacity. Currently there are a number of infrastructure projects running with the goal of building the required physical infrastructure for hosting cloud services. Probably the most notable is the construction of

an innovative container-based datacentre that will be installed in the banks along the Louros River in Central Greece, located near a hydroelectric dam operated by the public power corporation (PPC). The datacentre will use water from the river in order to cool the IT equipment and electricity produced by the dam, making it an excellent example of Green IT implementation. This datacentre will act mainly as disaster recovery of a second datacentre in Athens, which is also currently being expanded in order to accommodate a larger capacity of physical resources. These activities are also partially motivated by GRNET's commitment to provide cloud computing resources to Government agencies in Greece whenever required.

The introduction of ~okeanos GLOBAL on the other hand can act as a blueprint for other NRENs who wish to offer similar cloud computing services to their users. Other NRENs would be able to use this platform to provide virtual machine computing facilities to their own user base. GÉANT's support to clouds activities has been created to enable NRENs to share their expertise and bring innovative cloud services to the research and education community.

~okeanos continues to evolve and expand its capabilities. As part of the EGI federated cloud task force, focus has been placed on interoperability and integration aspects. This is primarily achieved with the native support for OpenStack APIs but also through the development of software integration layers that enable support for other standard APIs like OCCI and CIMI. GRNET realises that cloud computing services of this scale cannot remain isolated but have to be integrated through federation layers in order to attract a wider user base and be able to reach to a broader range of applications. In addition, innovative features, such as those developed by European CELAR project, can provide a significant technological advantage, help differentiate from the rest of the cloud offerings and compete head-to-head with larger commercial offerings in this area.

References

[1] GRNET, http://www.grnet.gr.

[2] ~okeanos, http://okeanos.grnet.gr.

[3] Synnefo cloud software stack, http://www.synnefo.org.

[4] Ganeti, https://code.google.com/p/ganeti/.

[5] EGI Federated Cloud Task Force, https://wiki.egi.eu/wiki/Fedcloud-tf:FederatedCloudsTaskForce.

[6] GÉANT, http://www.geant.net.

[7]~okeanos GLOBAL announcement, http://www.geant.net/MediaCentreEvents/news/Pages/New-Cloud-Computing-service.aspx.

[8] CELAR project, 7th Framework Programme, Software & Services, Cloud, DG CNECT, http://www.celarcloud.eu.

Enabling Swiss Researchers with Cloud

Dean Flanders, Friedrich Miescher Institute, Peter Kunszt, University of Zurich, Sergio Maffioletti, University of Zurich

There is a fundamental misconception that cloud is about technology. Cloud is actually a business model enabled by technology, and our focus is supporting researchers in Switzerland with cloud-enabled resources and tools to offer them "science as a service". This will help to reach the EU vision of a single European research area, and empower collaborations between institutions and across research domains. As well as allowing for efficient utilisation of resources and giving researchers more time to focus their energies in areas of innovation.

Who stands to benefit and how

Traditionally e-infrastructure topics have been targeted at national and institutional ICT providers. However, it is often the case that research groups have diverse needs that are beyond the scope of institutional or national level ICT providers. Instead, we must focus on working directly with researchers to cloud enable their resources or to use cloud resources (national or international), as well as to help build and support evolving research collaborations. This will allow researchers and technology providers to work together to achieve economies of scale and leverage, as well as re-use of existing investments. This is not just for large collaborations, but also small communities of researchers, which until now have been largely forgotten.

Addressing key concerns impeding the mainstream adoption of the cloud

The commercial world has moved in mass to cloud services because of their lower overall cost [1]. For a variety of reasons, the academic sector has been lagging behind in this move. It is critical this is addressed in order to increase the competitiveness of academic research. In Switzerland we ran a one year project that ended in April 2013 entitled "Academic Compute Cloud Provisioning and Usage" [2], which looked into cloud business models for HPC usage. One conclusion was that academic research is largely CAPEX driven, whereas cloud is heavily OPEX based, so this is one of the key reasons why cloud adoption has been lagging behind in academia. The European Commission has emphasised that, "as most knowledge creation and transfer uses digital means, all barriers preventing seamless online access to digital research services for collaboration, computing and accessing scientific information and to e-infrastructures must be removed by promoting a digital European Research Areas" [3]. The only viable way to achieve this is through the use of cloud resources, many of which will likely be from commercial providers who provide a rich set of resources across national borders.

Business models and sustainable services

In order to overcome obstacles and achieve this vision, the aim is to establish a national consortium for Swiss

e-Science Support (CHeSS), enabling collaboration of the local Science IT support organisations, allowing them to support projects of national and international scope, as well as to create synergies and national economies of scale. This has been done in co-operation with the Swiss National Grid Association (SwiNG) in order to build a Swiss eScience research infrastructure to support national and international eScience activities. In doing this, existing activities of consortium partners will be coordinated and where possible worked on co-operatively. In order to tackle the OPEX obstacle, models of co-operation with commercial providers are being explored where cost intensive resources may be run locally but by external companies to meet the needs of one or more institutions. Additionally, methods of innovation management such as crowd sourcing and crowd funding will be employed to find synergies and build co-operation. Solutions for researchers and institutions will be developed so they are re-usable by others, usually using SaaS approaches to ensure they are multi-tenant and achieve maximum efficiency. In particular it is important to collaborate as well at the international level to ensure that Switzerland has the widest range of resources within its digital ecosystem to ensure the competiveness of its researchers.

Links and References

[1] See for example, Novartis using Microsoft Office365 case, http://goo.gl/DLCN5e and Roche moving over 90,000 employees to Google Apps, http://goo.gl/jPJn8L.

[2] Results from "Academic Compute Cloud Provisioning and Usage" project presented on April 29th 2013 at the "Academic Compute Cloud Experience Workshop", http://goo.gl/mAlf6E.

[3] Brussels, 17.7.2012, COM(2012) 392 final, "A Reinforced European Research Area Partnership for Excellence and Growth", http://ec.europa.eu/euraxess/pdf/research_policies/era-communication_en.pdf.

Author affiliations: www.fmi.ch and www.uzh.ch

Scientific cloud computing using e-Science Central

Simon Woodman, Hugo Hiden and Paul Watson, Newcastle University

The team behind the e-Science Central [1] within Newcastle University has extensive experience in incorporating a wide variety of cloud computing technologies into research projects throughout our institution. Our users range from medical researchers analysing therapeutic data gathered directly from patients to groups of materials scientists sharing spectral data and analysis methods with colleagues distributed around the globe. Our involvement in these research projects has given us a unique insight into the issues both perceived and real faced by researchers when they try to transition their work from a small scale desktop computing environment to a complex distributed cloud platform.

Who benefits from e-Science Central

By combining our expertise in cloud technologies and the e-Science Central cloud platform, we enable other researchers who are looking to harness the power of cloud computing to build on top of a proven, open source system that addresses many of the concerns scientists and non-IT focussed users perceive when they are faced with making use of the cloud. Specifically, we can provide technology that presents a familiar user interface to complex cloud technologies, can scale out to hundreds of cloud server instances and can be used as a back-end data processing and analytics platform upon which to build custom applications.

Supporting wider uptake of cloud for research

Writing applications which are secure, reliable and scalable is an issue for any application developer. This is particularly challenging when academic scientists and developers from SMEs are domain experts in a non-IT field. These two areas, together with education and skills shortages are what we consider to be the greatest barriers to the large-scale adoption of cloud computing. However, this group, which can be seen as 'the long tail of science and industry', also has the most to benefit from the democratisation effect of cloud computing.

Middleware platforms are able to help solve some of the technical issues involved with transitioning scientists to the cloud. By structuring their application in certain ways, scientists are able to make use of a platform that gives them many of the capabilities of cloud computing without having to write all of the complex middleware themselves. e-Science Central provides services for secure, versioned data storage, archiving into cheap long term storage, data analytics through workflow technology which scientists can extend and enhance and an audit trail via extensive provenance capture. The e-Science Central analytics platform relies on workflows that can be scheduled across many cloud instances to provide scalability. Scientists can extend the in-built set of services with their own proprietary code written in a variety of languages. Further, they can also make use of a suite of workflows developed by colleagues and other users.

The use of workflows allows researchers to graphically define complex data processing pipelines. The various operations in these pipelines can be written in a range of programming languages that are specifically targeted at scientists and engineers, for example, the R statistical framework and the Octave language. Once created, these processing pipelines can be applied in parallel to any number of data sets and shared amongst other researchers to build a library of reusable best practice protocols.

Through the use of seminars, training courses and online resources we are trying to educate scientists in the benefits of cloud computing. Many do not realise the existence of such pay as you go resources and would struggle to justify the cost of large dedicated resources. Groups such as the Digital Institute at Newcastle are able to collaborate on grants to help the migration of both the data and analysis code lowering the time costs to the scientist.

Links and References

[1] http://www.esciencecentral.co.uk.

EMBL-EBI's Embassy Cloud: Bringing computation to large data sets

Steven Newhouse, Andy Cafferkey, Ewan Birney, The European Molecular Biology Laboratory

The European Bioinformatics Institute [1] in Cambridge UK is an outstation of the intergovernmental European Molecular Biology Laboratory (EMBL-EBI) providing free available data from life science experiments. EMBL-EBI has to date just provided web-based access (both interactive and programmatic) to its tools and data. Researchers wishing to undertake more extensive and intensive analysis would need to download the relevant data sets to their local resources, duplicate any dependent services and perform their data analysis locally. EBI has, over the last two years, through the 'Embassy Cloud' been developing a cloud based model to data analysis that is changing the way computational biology is being undertaken.

Who stands to benefit and how

Given the continued exponential growth in public data sets the ability to download and establish a local analysis environment is always demanding greater resources and expertise. EMBL-EBI's Embassy Cloud provides a secure environment where 'tenants' have access to virtualised compute and storage resources under their own control and management, while being able to have direct access to EMBL-EBI's hosted data sets – bring their analysis activities to EMBL-EBI's hosted data sets. Thus, computational biology researchers benefit by being able to focus on their research activities rather than becoming infrastructure experts by having close network access to data and services.

Sustainable big data and cloud services

EMBL-EBI's Embassy Cloud brings forward a paradigm where compute moves to the 'big data' as opposed to the model that has dominated to date of bringing remote data to the local compute. This paradigm shift is being brought about by the exponential growth in the public data sets being seen in the life-sciences community and the growth of clinical data sets which have tight regulatory supervision and control. In order to unleash the benefits that personalised medicine and improved diagnosis tools can bring, complex analyses that bring together confidential medical data and public 'big data' sets will need to become routine and reliable from both a technical and regulatory perspective.

To establish such an operating model, all stakeholders around the clinical data sets will have to have confidence in the underlying technical infrastructure. At present, offering access to defined resources within the same legislative domain as the data was collected, provides a measure of confidence. The 'Embassy' model brings together the flexibility of on-demand provisioning with the confidence that a secure bounded set of resources can bring within a multi-tenanted environment.

Sustainability of this model through its transfer to the commercial sector has two main aspects: the willingness and expertise for commercial providers to replicate the public data sets, and the ability to

provide a trustworthy environment that is aligned with the data to which access is managed. This is an example of 'Information as a Service' where the hosted data provides a 'centre of gravity' that attracts compute work and potentially supports the cost of providing the environment.

While the initial implementation of the Embassy Cloud has focused around its provision at EMBL-EBI in a single controlled environment, it is a model that could be replicated at other sites within the biomedical research infrastructures (such as ELXIR, [2]) if certain concerns can be addressed:

» Data Confidentiality: How can data sets containing confidential information be moved to an arbitrary cloud, accessed by just the authorised VMs, and be removed with confidence from the remote cloud environment?
» Reproducible and Reliable Analysis: How can a researcher or clinician have confidence that the analysis framework that they are using is operating correctly and has not been altered since its release? How can this analysis be recorded so that it can be reproduced or audited at a later date?
» Data Movement and Access: How can large potentially sensitive data sets be moved between legal jurisdictions so that the user is able to make the most flexible use of the resources available while complying with the restrictions imposed by the data owner?

Mitigations can be established for many of these issues, however these mitigation impose constraints and flexibility on the cloud resources that can be brought to bear on a particular analysis, and long-term these need to be removed.

Links and References

[1] http://www.ebi.ac.uk/.
[2] http://www.elixir-europe.org/.



EU Brazil Cloud Connect – Addressing societal challenges in the cloud

Ignacio Blanquer, Valencia University of Technology, Francisco Brasileiro, Federal University of Campina Grande

EU Brazil Cloud Connect [1] is a new international co-operation project aimed at accelerating scientific discovery to advance knowledge on several challenges of high social impact. By creating a joint cloud infrastructure, EU Brazil Cloud Connect will enable a multi-disciplinary user community to co-operate across borders at different levels: infrastructure, federation, programming models, applications, usage by research and public institutions, and ultimately, wider uptake of sustainable services and tools.

EU Brazil Cloud Connect is aimed at driving advances in three key areas:

» Virtualised resource federation using clouds that promote sustainable services and tools.

» Programming frameworks in the cloud, including big data analysis.

» Requirements of scientific applications demanding high capability computing and data in the cloud, especially **epidemiology**, **heart simulation** and **climate change**.

By building on the success stories of earlier initiatives like VENUS-C, EUBrazilOpenBio and MyScientific Cloud [2], EU Brazil Cloud Connect will extend integrate functionalities that ensure applications effectively scale across federated infrastructures. This approach will improve not only the experiences of developers using tools to deploy SaaS systems but also greatly facilitate researchers using them to tackle grand global challenges.

Who stands to benefit and how

EU Brazil Cloud Connect is aimed at demonstrating the efficiency and cost effectiveness of tools and solutions designed to address societal challenges of common interest to Brazil and Europe.

Leishmaniasis Virtual Laboratory - *Anticipating outbreaks.* Every year 1-2 million new cases of Leishmaniasis occur. More effective control of neglected tropical diseases like Leishmaniasis is vital to achieving poverty reduction and spurring social-economic development without waiting for countries to fully develop and living conditions to improve over a potentially long period of time [3]. Cloud Connect will tackle this challenge by improving the knowledge on the susceptibility of the outburst of the Leishmaniasis disease by integrating species distributed databases of parasite vectors and biomolecular data from parasites with bioinformatics and Niche Modelling processing pipelines.

Whole Vascular System Simulation - *A Pureblood simulator.* Cardiovascular diseases have a huge impact on population, particularly people with a low- and middle-income [4]. Simulating a heartbeat is a complex, multi-scale problem. Cloud Connect aims to achieve an important advance in the exploitation of high-level heart simulation data in the context of the Virtual Physiological Human. To achieve this goal, it will deploy a complete blood simulation system with accuracy beyond the state of the art by integrating the heart simulation system (ALYA) with a complete vascular simulation system (ADAN).

Biodiversity & Climate Change - *Breaking the vicious circle.* It is vital to understand the mutual interaction at a global scale between climate change & biodiversity dynamics. Cloud Connect is focused on generating new knowledge on cross-relations between climate change and terrestrial biodiversity through the use of earth observation and ground level data together with simulated data. The project will integrate two workflows combining models of plant species distribution and multi-level imaging data and processing in a scientific gateway.

Open standards and interoperability

Cloud Connect brings expertise and practical experiences from both enterprise and research on standardisation initiatives, ranging from contributions to the EGI Federated Cloud [5] and interoperability testing to co-authoring the SIENA Roadmap aimed at building consensus around standardisation initiatives globally [6]. By implementing relevant standards that have already proven effective in similar contexts and performing interoperability testing through the Cloud Plugfests series, Cloud Connect will avoid duplication of efforts and contribute to standardisation initiatives globally. A good case in point is OGF's Open Cloud Computing Interface (OCCI), which is emerging to be the dominant interoperability standard of choice for the federated control of multiple cloud provider infrastructure interfaces.

Cloud Connect will also coordinate interoperability with Helix Nebula [7]. Cloud Connect will provide technical expertise to support the porting of applications and funds for Helix Nebula resource consumption. The aim is to evaluate cloud for science in different contexts while establishing an important link to the vision for an e-infrastructure commons marketplace [8].

Last but certainly not least, Cloud Connect is taking the Cloudscape series to Brazil, with two workshops already planned for 2014 and 2015. Key goals include:

» Educating communities on the benefits of interoperability, portability and open APIs, whether that be to avoid vendor lock-in and high migration costs, broaden choice or enable the more efficient allocation of resources. Looking into increasing consumer concerns around these issues and who should do what to solve them.

» Sharing experiences in cloud deployment and usage, exploring opportunities for wider uptake of Cloud Connect by supporting requirements identified through project synergies. This includes investigating success factors and metrics to gauge impact.

» Exploring opportunities for future R&D, including the involvement of start-ups and businesses with EU-Brazil mobility programmes and funding opportunities in mind.

Business models and sustainability

Cloud Connect has already identified 16 coarse-grain potentially exploitable assets that will ensure long-term sustainability through dedicated business plans. Examples of exploitable assets include improved and more efficient services for collaborative research; commercially exploitable assets, assets for higher education institutions and public authorities. This initial identification of assets will lay the foundation for activities performed during the project life-cycle. Specifically, the project will analyse the demand and supply sides, position its assets in the market, and identify new opportunities for collaboration between

research-research; research-public authorities and research-industry, including the potential for new spin-outs for data services and public-private partnerships for service provision. The outcomes of these analyses will help define business models and exploitation plans.

Links and References

[1] www.eubrazilcloudconnect.eu.

[2] www.venus-c.eu; www.eubrazilopenbio.eu; (in Portuguese) https://www.lncc.br/sinapad/projectmanager/public/projects/gt-mcc, http://www.rnp.br/pd/gt2012/GT-MCC.html.

[3]. World Health Organisation, *First WHO Report on neglected tropical diseases. "Working to overcome the global impact of neglected tropical diseases"*, 2010, http://www.who.int/neglected_diseases/2010report/en/.

[4] I. LeGrice, P. Hunter, A. Young and B. Smaill, 'The architecture of the heart: a data based model', Phil. Tans. R. Soc. Lond. 359, 1217-1232, 2001.

[5] EGI Federated Cloud, https://www.egi.eu/infrastructure/cloud/; on technical activities, see https://wiki.egi.eu/wiki/Fedcloud-tf:Main. See also, M. Drescher, 'Interoperability is the key to freedom in the Cloud', *Cloudscape VI Position Papers*, 2014.

[6] The SIENA Roadmap on Distributed Computing Infrastructure for e-Science and Beyond in Europe, June 2012, http://www.cloudscapeseries.eu/Content/CloudscapeUpdates.aspx?id=96.

[7] http://helix-nebula.eu/.

[8] M. Lengert, B. Jones, D. Foster, S. Newhouse, 'e-Infrastructure Commons Marketplace', *Cloudscape VI Position Papers*, 2014. See also, http://www.e-irg.eu/images/stories/dissemination/white-paper_2013.pdf.

Funding

EU Brazil Cloud Connect is funded under the European Commission's 7th Framework Programme for Cooperation and the Brazilian CNPq (Conselho Nacional de Desenvolvimento Científico e Tecnológico), Ministry of Science and Technology.

Boosting business innovation in the Cloud

Flexibility in financial services through the cloud

Stephen Watling, Deutsche Börse Group, Robert Jenkins - CloudSigma

This paper elaborates on the work between CloudSigma [1] - an innovative public infrastructure as a service provider, and Deutsche Börse [2] - one of the largest financial exchange organisations worldwide in deploying a proprietary cloud-based solution to provide trading members with self-provisioned, on-demand access to Eurex' s T7 trading architecture for testing and development purposes.
This paper not only covers the purchasing requirements and considerations but also the approach and challenges of a large corporate operation looking to leverage public cloud infrastructure and the experiences of launching a cloud-based application in an organisation more accustomed to traditional deployment on dedicated infrastructure.

Who stands to benefit and how

Both end users and service providers will gain insight into the approach and process of this real-world use case resulting in a successful deployment of a new product by a large corporate entity using public cloud. Deutsche Börse's experience in leveraging public cloud is directly applicable to any larger organisations with strict data security requirements currently considering incorporating public cloud infrastructure into their future strategy. Key success criteria and requirements are outlined, as well as lessons learned and best practices as a result of both organisations' experiences to date.

Business models and sustainable services – Meeting requirements

Deutsche Börse did not start out to build an application in the cloud when they started work on Virtualised Private Simulation. On the contrary, the project team tried to avoid the cloud approach. They feared resistance based upon the many misconceptions that are often associated with cloud technology. However, upon weighting the requirements of the project, including low cost to implement (with minimal CAPEX), fast development cycle, highly varying usage patterns and high availability from many geographic locations, the facts were more than enough to overcome any resistance within the organization. The only practical option was to launch Virtualised Private Simulation in the cloud.
In order to build expertise in deploying cloud-based applications quickly, the Deutsche Börse VPS project group worked with CohesiveFT, a firm that specialises in cloud-based application deployment. With CohesiveFT's assistance they analysed requirements and evaluated cloud providers to identify the best fit.
The technical requirements are as follows:

» Very fast virtual machine start-up time.
» High degree of data privacy (ensuring adherence to German privacy laws as well as Deutsche Börse Group data security policies).
» Compatibility with existing proprietary software without modifications.

» High-capacity virtual machines (12-18 cores, 24 – 32 GB RAM).
» Support for user-uploaded disk images of >50GB size.

As such it was critical that the chosen cloud matches closely the existing production environments of Deutsche Börse's trading systems whilst at the same time is offering stability, performance and a pricing model that made sense for an on-demand SaaS platform.

They finally settled on working with CloudSigma in their Zurich data centre. CloudSigma provided the reliability, capacity and flexibility required by Deutsche Börse at a competitive rate.

Addressing key concerns impeding the mainstream adoption of the cloud

Project teams often face resistance when considering cloud-deployment for mission-critical or highly proprietary systems. Deutsche Börse's VPS development team was no different. In order to address key stakeholders' concerns, the project team spent a lot of time working with them to identify their pain points and took these challenges into account while designing and implement VPS.

The two biggest concerns in deploying VPS into the cloud for Deutsche Börse were data privacy & security and operational impact.

Data privacy & security: As a German company, Deutsche Börse must comply with German privacy laws as well as strict company policies on data security. To address the privacy concerns, the VPS team designed the system so that no user information is stored in the cloud. To ensure data security, the project team worked with group security experts to build encryption into every aspect of the system. Lastly, by choosing a Swiss-based cloud provider, Deutsche Börse also started to benefit from the strict data and security protections in Swiss law.

Operational impact: Security and privacy of course are important. But so are reliability and supportability. Identifying a cloud solution that provides the reliability their customers expect was important for the VPS team.

Lastly, it was important to develop new support tools and to update operational policies and procedures to work in the cloud. Existing tools and support procedures were not sufficient to extend into the cloud.

CloudSigma proved to offer the best combination of reliability and supportability balanced with a high degree of privacy and data security protections.

Links and References

[1] http://www.cloudsigma.com/.

[2] http://deutsche-boerse.com/dbg/dispatch/en/kir/dbg_nav/home.

Startups and small businesses in the cloud – Experiences from Cloud Software Finland

Janne Järvinen, F-Secure Corporation

F-Secure has recently finished leading a large (Đ60 million funding with 30 organisations), four-year (2010-2013) Finnish Cloud Software Programme that has built competencies and capabilities to be successful in the cloud. The total gain to the participating companies has been estimated to be as much as Đ100 million a year with 100s of millions in new business potential. In the programme, open cloud platforms, lean & agile ways of working and changes to business operation models were studied comprehensively. The creation of a superior user experience and safe cloud solutions was a common theme across the research. The core results of the programme have been published in several guidebooks, which are electronically available at on the website of the Cloud Software Programme [1].

Our Quick Guide to Cloud Success

One of the Cloud Software guidebooks provides key learnings from Cloud Software Finland for businesses to plan their journey to the cloud [2]. The aim was to create a concise guidebook containing practical, evidence-based information and descriptive guidance for organisations wishing to be successful in the cloud. This will help unleash the potential of Cloud Computing in Europe. This guide includes insights, experiences and examples from Cloud Software Finland, the European Cloud Partnership and other interested parties. This short guide has been designed to offer a thoughtful read for anyone interested in starting to use cloud or enhancing their use of cloud.

The Cloud Software Program's "Quick Guide to Cloud Success" captures key experiences from the Digile SHOK Cloud Software. This informative, easy-to-read book presents current methods with examples on how best to move to the Cloud by showcasing Finnish organisations and their experiences. In addition, there are many helpful tips on how to embrace cloud or use cloud-based services and offerings. Many cloud solutions and tools are easy to use, free or competitively priced and available to use by anyone. For example, a low cost Cloud-based service can be set up in minutes.

How to plan your journey to the cloud

Cloud technologies are facilitating an on-going wide-scale transition towards the digital economy, which is emerging rapidly and impacting all of us. For example, the recent revelations of international web espionage have accelerated the need to set up datacentres in Europe. This move offers fantastic opportunities for Finnish software companies offering cloud services.

However, technology alone is not the only aspect to be taken into account when moving to the cloud. Cloud technologies do offer new opportunities for companies to be competitive, but technology does not solve everything. Too often, companies make a decision to move to Cloud because of technology alone.

It is important to note that a move to Cloud often means working practices and business models have to radically change in order for them to succeed against the competition.
We highly recommend setting in motion the following transformations, with a special emphasis on user experience, security and sustainability: cloud business, cloud technologies, lean and agile organisations.
The "Quick Guide to Cloud Success" presents examples that can be used by both companies and public sector administration. In addition to adopting technology, the Guide emphasises a business activity-centred approach and incorporates user experience, data security and agile methods in developing working practices. Based on the experiences of companies that took part in the Cloud Software Program, the Guide highlights their key learnings as the companies that have developed and launched cloud services for global markets. The Guide provides a perspective of five stages:
Stage 1. Planning and investment.
Stage 2. Implementation.
Stage 3. Checking and metrics.
Stage 4. Preparing for the worst.
Stage 5. The future.

Links and References

[1] Cloud Software Programme, www.cloudsoftwareprogram.org.

[2] The Quick Guide to Success: Key Learnings from Cloud Software Finland, https://www.cloudsoftwareprogram.org/results/deliverables-and-other-reports/i/29042/1941/quick-guide-to-cloud-success, produced in collaboration with the European Cloud Partnership –consortium, www.ec.europa.eu/digital-agenda/en/european-cloud-partnership.

Funding

The Cloud Software Program (2010-2013) was a research and development effort funded by Tekes, the Finnish Funding Agency for Technology and Innovation. The partnership included over 30 companies and research institutions, working together over a 4-year period.

MobiCloud - a novel cloud-based platform for cross-platform context-aware enterprise mobile apps

Vladimir Bataev, EsperantoXL, Xavier Aubry, Appear

MobiCloud [1] enables companies to quickly mobilise existing line-of-business applications that sit in their back offices, and reducing integration efforts in the process. The MobiCloud community is open to new members: mobile developers, system integrators and enterprises. This award-winning consortium [2] is made up of five companies from four countries, co-funded under the European Commission's Competitiveness and Innovation Programme (CIP) focused on building a new generation cloud-based platform to create, manage and deploy cross-platform mobile apps for enterprises.

MobiCloud

MobiCloud enables the roll-out of industry app stores, where multiple apps could be bought and sold without the hassles associated with the more traditional app stores like the ones from Apple or Google. Company-specific app stores are possible as well, which opens the door for creation and delivery of completely new services to the companies by both internal and 3rd party developers.

How we do it. MobiCloud applications are HTML5 apps that run in a native device container, which makes them easily portable across different platforms while giving access to device hardware. Additional work is focused on quickly moving elements of these legacy applications to the cloud.

Who stands to benefit and how

Multiple parties stand to benefit from using the MobiCloud platform, especially small and medium-sized businesses, looking to boost their position in the marketplace. MobiCloud enables companies to offer employees mobile access to either existing software or completely new mobile solutions. As end-users of mobile apps, employees can use their own or company-provided devices to benefit from the availability of critical data or business processes 'anytime, anywhere'.

Developers also benefit from MobiCloud, because it is easier to develop quality mobile apps. The MobiCloud platform provides context-aware capabilities out of the box that allow developers to create innovative apps that take into account the specific requirements of users.

Business models and sustainable services

MobiCloud is creating a commercial technology that is fit for use by companies by addressing their specific pain points. Based on four trial cases in rail, light city transit, construction and field services, MobiCloud is showcasing a set of individual mobile apps running on the platform and demonstrating the viability of commercialisation.

Behind both our pilots and the platform, there is a proven business model generated by a careful iterative customer delivery process. Two of our trials, construction at London Bridge by Costain UK and the Swedish transport company Tågkompaniet, are already run on a commercial basis far ahead of our original schedule. MobiCloud is also creating a marketplace by attracting other customers, including developers who want to use the platform apps either for companies that already run trial case solutions and companies who want to have new apps created for them.

[1] http://www.mobicloudproject.eu/.

[2]http://ec.europa.eu/digital-agenda/en/blog/eu-funded-project-mobicloud-wins-multiple-awards-its-mobile-cloud-solution-construction.

Transmetrics – A Cloud solution that brings big data for cargo transport

Asparuh Koev, Transmetrics

Transmetrics brings Big Data to cargo transport (a Ð2 trillion industry) to improve capacity utilisation. This industry spends over half of its resources on "transport air", that is, empty spaces with nothing inside. Running cargo vehicles in empty or even half-full spaces has staggering economic costs. What's more, it produces cargo emissions, burns fossil fuels and creates traffic jams.

At Transmetrics, we build a cloud predictive analytics product, which merges external signals with historical shipping data of transport companies. The result is a prediction of future shipping volumes 3-6 weeks ahead of time, giving transporters a chance to eliminate empty spaces before they happen.

Who stands to benefit and how

Cargo transport companies will benefit from our cloud solution by significantly reducing the number of vehicles it takes to transport a given volume of shipments, which uses capacity more efficiently. Our data-driven simulations show that a transport company's current profit margin can be multiplied 3x to 6x.

The general public will benefit by lowering CO2 emissions, and by having less traffic on the roads. Also, the build-up of petabytes of detailed transport data in a central location will ultimately enable a more intelligent understanding of the dynamics of transport, and may open up other economic opportunities.

Addressing new challenges on the horizon

Our society is experiencing a once-in-a-generation shift in consumer patterns. For the first time this Christmas, shoppers predominantly purchased on the Internet, rather than in department stores. This is leading to capacity problems at large transporters, such as FedEx and UPS, who were not able to deliver hundreds of thousands of shipments on time. At the same time, most traditional retail stores reported disappointing results.

This trend will only continue in the next decade, as e-commerce matures, trust in it grows, and e-commerce players like Amazon.com start entering more and more the business domain of retailers such as Wall-mart.

In the cargo transport industry, this will lead to a tectonic shift in operations. Where, in the past, the main shipping unit was the container or the lower level, the pallet, now many more shipments will have to be transported in smaller sizes, boxes, even envelopes. Handling such smaller packages will require transport networks that are much more rigid than today. We see even mid-sized players starting to gravitate towards groupage and fixed departure networks for small shipments.

The more fixed the network is, and the more it deviates from the container size, the more empty space it has. Fixed, small package networks traditionally have had very low capacity utilisation. As an example, a container transport operation is typically 24% empty kilometers, while an average groupage network is on

average 43% empty, and for parcel networks the empty space could go to 50-60%.
Therefore, the shift in consumer behaviour will lead to ever decreasing levels of efficiencies for transport companies. Without proactive action, the outcome will be even more trucks on the road, and higher prices for transport, which will contribute to slowing down economic growth.
This shift is creating much interest within the industry into how big data can be used to predict several weeks ahead of time the volumes of smaller parcels. That will enable a transport company to counter-act the emerging drivers of inefficiency, and even improve on today's situation, by proactively adjusting their transport network to have just the right network capacity on a given day. Our simulation shows that for groupage network, prediction enables the empty space to be reduced from 43% down to 18%. That means a decrease of about 20% in the amount of vehicles travelling, and leads to significantly improved profitability for the transport company.

Addressing key concerns impeding the mainstream adoption of the cloud

Today, transport companies run on legacy transport management software (TMS). These are often terminal-based mainframes, with limited computing and data storage. As these systems are very complicated, and intimately tied to daily operations, the migration to cloud-driven TMS has not even started yet, and will take decades to accomplish.
We offer the alternate path of "liberating" the data that today sits inside these systems, so that while the TMS remains local and limited, the data it generates can be easily moved to the cloud, stored in a generic industry format, and analysed to achieve business benefits. In addition to the benefits for any company contributing data, the accumulation of data from many companies in a central big data infrastructure will, for the first time in history, create a detailed, precise description of shipping flows globally.
Consider this. The largest transport company (DHL) has just 3% of the world's transport volumes, and it goes down pretty fast from there. With the data being hidden and fragmented, no one, not even the biggest market players, has detailed data on how the transport flows really work. By combining data from thousands of companies, we will achieve this for the first time, for the benefit of companies, researchers and humankind.

Links and References

[1] http://transmetrics.eu/.

Scalable Data Analytics – A new start-up in the cloud

Domenico Talia, University of Calabria & National Research Council (CNR)

Scalable Data Analytics [1] is a start-up founded by a team of researchers with a background in computer engineering from the University of Calabria and the Institute of High Performance Computing and Networking within the Italian National Research Council (ICAR-CNR).
This award-winning start-up [2] develops an innovative, high-performance Software-as-a-Service (SaaS) system, which is able to efficiently analyse large amounts of data and information by exploiting Cloud computing technologies. The main innovation behind the system is an efficient and scalable algorithm for parallelising data analysis applications modelled as complex workflows. The algorithm efficiently exploits the vast storage and computing potentialities of Cloud systems.

Addressing new challenges for data analytics

Pervasive and huge digital data repositories are increasingly becoming a big part of our daily lives. Data warehouses, web pages, streams, tweets and posts are making digital data bigger, more complex, and ubiquitous. This situation requires smart techniques for data analysis and scalable architectures to enable the efficient extraction of useful information and knowledge from data. Big data and cloud computing are natural allies, so research work in this area must enable the shift from a computation and data management infrastructure to a pervasive and scalable data analytics platform. This trend needs new models and technologies that enable cloud computing systems to support the implementation of clever data analysis algorithms that are scalable and dynamic in resource usage on the cloud.
Complex data mining and knowledge discovery tasks involve data- and compute-intensive algorithms, which require large and efficient storage facilities together with high performance processors to get results in acceptable times. Cloud computing infrastructures can play the role of an effective platform to address both the computational and data storage needs of big data mining applications. What's more, we need new solutions for implementing cloud-based data analytics services, programming tools, and applications. Scalable Data Analytics looks at the market of companies that hold big data. Much of this data already resides in the cloud, and this trend will increase in the future. Tackling and gaining value from cloud-based big data is the mission of Scalable Data Analytics.
Over the next few years, cloud-based data analytics clouds are expected to become common platforms for big data analytics. Both the Platform as a Service (PaaS) and Software as a Service (SaaS) models can be adopted for implementing big data analytics solutions on Clouds. PaaS can support data analytics programming suites and environments where data mining developers can design scalable data analytics services and applications. The SaaS model offers complete big data analytics applications to end users that can execute analysis on large and/or complex data sets by exploiting the scalability of Clouds both in data storage and processing power.

Who stands to benefit and how

Big data analytics is the advanced use of mining techniques on very large and complex data sets. Whether for research or business, data analytics techniques and tools help people to dig data and extract information and knowledge useful for making new discoveries or for making smart decisions that improve the business process or service. In other words, putting big data and knowledge discovery techniques together with scalable computing systems, like cloud computing, produces new insights faster. But very few cloud-based analytic platforms are available today despite known benefits for both public and private organisations.

Scalable Data Analysis looks at the market of companies that hold Big Data as a result of their business. Much of this data already resides in the Cloud, and this trend will increase in the future. We have recently implemented a Cloud Data Mining Framework as a high-level PaaS data analytics programming environment, which also provides a set of SaaS suites for big data analytics built on the PaaS layer. This approach enables end users to make complex analysis without having to know details of the platform layer or how the analytics suite has been programmed.

The data analysis framework is a service-oriented workflow-based on a software environment for designing and running big data analysis applications on cloud platforms. This framework provides a workflow programming interface for running service-oriented applications on a cloud infrastructure. In this environment, developers can combine data sets, analysis tools, data mining algorithms and knowledge models that are implemented as single web services. The workflow paradigm is exploited to compose all these services in distributed workflows and execute them concurrently on different virtual machines.

Developers can use workflows, which consist of complex graphs of many concurrent tasks, to address the complexity of business and scientific data analysis applications. This approach supports data analytics design by providing a paradigm that encompasses all the steps of data analysis, from data access and filtering to data mining and interpreting the knowledge generated. The system provides both visual and script-based workflow programming, so as to meet the needs of both high-level users and skilled programmers. In addition, the system is open to third-party tools for easily importing existing algorithms into the programming environment, such as custom data mining tools or algorithms from open-source projects.

Links and References

[1] http://scalabledataanalytics.com/.

[2] http://scalabledataanalytics.com/news/.

ClouDesire – a new cloud based app store for software vendors

Eddy Fioretti

ClouDesire [1] is a start-up company offering innovative ways of on-line selling and the distribution of applications over the web founded by a team from Milan Polytechnic and the Scuola Normale Superiore of Pisa. It is based on a cloud platform which enables the easy management of commercial activities or the distribution of applications without burdensome investments. More broadly, it contributes to socio-economic innovation by promoting the growth of the digital economy across small businesses.

Addressing new challenges

In order to succeed in a highly competitive global market place, SMEs need to embrace technologies that can speed up time to market. Expensive investments, mission-critical and time-consuming tasks all too often stand in the way of innovation and competitive edge. Our analysis of the cloud landscape from the perspective of a small- to medium-sized software vendor shows the following challenges impeding their transition to the cloud:

» Acquiring the required cloud-specific know-how.
» Allocating human resources for software re-engineering.
» Profoundly changing marketing and sales processes.
» Adapting post sales and maintenance support to customers.

Similarly, service providers have several valuable competitive assets (e.g. digital infrastructure, high quality media/services) and a large and loyal customer base that they'd like to leverage. These providers face the challenge of what to sell and how to deliver to a still unexplored target market of small businesses.

ClouDesire is committed to solving the challenges software vendors face in today's global marketplace by providing ready-to-sell SaaS apps to service providers, bridging the gap across infrastructure, applications, and users. ClouDesire fits with every Cloud Provider and all applications are welcome, irrespective of the development language used.

Who stands to benefit and how

ClouDesire enables software vendors to sell and distribute their apps as SaaS by automating several expensive, mission-critical and time-consuming tasks, and end-users to find them, thus overcoming entry barriers, especially for small businesses which help to drive the European economy.

ClouDesire brings the following benefits:

» Fostering the full take up of the cloud in the market of small and medium-sized software companies.
» Easing the proliferation of SaaS applications produced by small and medium-sized software vendors.
» Boosting the marketing & sales opportunities of small and medium-sized software vendors.

» Offering a concrete solution to the changing business models of small and medium-sized software vendors in times of economic difficulties.

Specifically, companies can target a worldwide customer base through multiple sales channels: their own app store, the ClouDesire marketplace or partnered service providers. They can also boost time to market and stay focused on their core work.

Links and References

[1] http://www.cloudesire.com/



CoherentPaaS: Coherent and Rich PaaS with a Common Programming Model

Ricardo Jimenez-Peris, Madrid Polytechnic

Addressing key challenges

Today, an application developer using multiple cloud data stores and SQL databases faces two main difficulties. The first challenge is that updates across data stores do not cater for loss of data coherence in the advent of failure (they are not "atomic"). The second challenge comes from the infeasibility of doing queries across data stores, since only APIs and/or query languages are provided for each individual data store.

The goal of CoherentPaaS [1], a project funded under the 7th Framework Programme, Software & Services, Cloud, DG CNECT, is to provide an integral platform for cloud data management. The platform will integrate a wide range of cloud data management technologies, including 3 NoSQL data stores, 3 database technologies and complex event processing, or CEP for short. The approach adopted by CoherentPaaS will overcome challenges of using multiple cloud data stores and traditional databases in a single application.

Who stands to benefit and how

Cloud application developers and Platform as a Service providers will benefit the most from the outputs and assets of CoherentPaaS. Key benefits range from full data coherence to an easy and common way of accessing data from different data sources: NOSQL data stores, relational databases and CEP.

The key competitive advantages of the platform are:

» Providing access to an arbitrary set of the data stores with full ACID guarantees across all the data stores, where ACID stands for atomicity, consistency, isolation and durability.

» Offering a common query language to make queries across multiple data stores.

Business models and sustainable services

CoherentPaaS will commercialise some of its main assets through a spin-off, which will be created by the coordinator together with some of the partners from the consortium. In addition, partners from enterprise will jointly commercialise other assets through the new spin-off venture. The goal is therefore to drive a success story that transfers European research excellence through uptake by enterprises while strengthening European enterprises and their product/service portfolio through research.

Links and References

[1] www.coherentpaas.eu.

Funding

CoherentPaaS is funded under the European Commission's 7th European Framework, Software & Services, Cloud, DG CNECT (FP7 ICT Call 10). This 3-year project started in October 2013.



StackSync: open source personal cloud for organisations

Pedro García-López and Ivan Utgé-Hernández, Universitat Rovira i Virgili

Addressing key challenges

Users will increasingly access their data from a variety of devices, operating systems and applications. Organisations will need to deal with a growing amount of data and take care that their sensitive data is not compromised as it will be no longer on their users' machines. The Personal Cloud model defines a ubiquitous storage facility enabling the unified and location agnostic access to information flows from any device or application. But Personal Clouds are in their infancy and two major problems must be solved: privacy and interoperability.

CloudSpaces, a European project funded under the European Commission's 7th Framework Programme, Software & Services, Cloud Computing [1], advocates a paradigm shift from application-centric to person-centric models where users will retake the control of their information. StackSync aims to become the next generation of open Personal Clouds, dealing with interoperability (avoiding vendor lock-in) and privacy issues. It will also take care of scalable data management of heterogeneous storage resources and will provide a high-level service infrastructure for third party applications that can benefit from the Personal Cloud model.

After over a year of focused development, we are now poised for the first launch of StackSync to the market of higher education and research institutions. To support this launch, we are collaborating with RedIRIS, a community of universities, research centres and government institutions in Spain. The main driver behind the interest of RedIRIS is our security model and the potential to reduce IT costs with our cloud storage technology. With regards to security, our solution provides zero knowledge for the cloud provider thanks to encryption technologies. All information sent to the cloud (metadata and contents) is then protected by symmetric keys only known by the user. Our security model is modular and configurable, so that users can configure specific folders to be encrypted or not depending on their privacy requirements. In this case, our system provides client-side encryption using AES-256 where the key is only known to the user.

In addition, StackSync will provide a privacy-aware sharing component that can work on top of existing cloud platforms (e.g. as a third party app) or can be integrated within existing cloud infrastructures. The goal of this component is to give the end-user indicators about the risk that is posed by sharing particular data items in particular contexts, and also propose and implement policies that they can adopt for mitigating that risk.

Interoperable solution based on open standards implementation

StackSync, the Open Source Personal Cloud specially designed for organisations, addresses the real needs of organisations and provides a cloud storage solution with scalability, openness, security and privacy awareness. Further, StackSync fits perfectly into any kind of organisation (SME, large corporations,

government, educational, etc.) offering adaptable cloud storage: private, hybrid or public.

Unlike home solutions, StackSync is using OpenStack Swift Object Storage technology, ensuring advanced data redundancy and scalability. StackSync is also based on an advanced synchronisation technology, similar to Dropbox, with data optimisation (chunking, compression, bundling and push mechanisms) that allows it to scale to thousands of users with an efficient use of cloud resources.

StackSync lets different Personal Clouds share and access information located in different cloud providers by using our Open APIs. This avoids vendor lock-in and facilitates information sharing between different services. Furthermore, StackSync is open to the community as an Open Source project, which also means that third-party applications can use our APIs to integrate their services on top of StackSync.
In the context of the FP7 Cloudspaces project, Canonical Ltd is integrating UbuntuOne with StackSync. It means that StackSync will benefit from feedback from the UbuntuOne community with millions of users. This would be the first real interoperability scenario between different Personal Clouds.

Who stands to benefit and how

StackSync is based on advanced synchronisation technology with data optimisation that allows it to scale to thousands of users with an efficient use of cloud resources. It means less IT infrastructures costs for organisations, compared to other competitor's solutions.
Infrastructure providers, Software providers, SMEs and Public Institutions are the four groups that will benefit from the StackSync project.
IaaS providers: StackSync helps IaaS providers to deploy an enterprise Personal Cloud for thousands of users at low cost (based on OpenStack Swift). They will be able to improve their service portfolio and engage customers with a Personal Cloud specially designed around their individual organisations.
Software providers: can take advantage of cloud storage features and bundle their software with a Personal Cloud (StackSync) focused on the organisation's specific needs in terms of security, scalability and openness.
SMEs: StackSync proposes two deployment scenarios focused on SMEs with the aim of boosting their storage in the cloud. StackSync-P, a private cloud that can be deployed on premise, and StackSync-H, which is a hybrid cloud that keeps metadata on premise and stores encrypted raw data on public clouds.
Public Institutions: StackSync keeps citizen data in a cloud storage platform that enables public institutions to manage large amounts of information without the risk of data security breaches while avoiding large investments on IT infrastructure.
Typically, academic institutions cannot afford to migrate their TBs of data to the cloud, because of the cloud provider costs. Instead of moving all data to the cloud, StackSync provides them hybrid or private cloud solutions that ease this big data migration, and also reduce their IT infrastructure budget by leveraging the efficiency of OpenStack cloud infrastructure.

In summary, StackSync provides organisations with an innovative Personal Cloud specially designed for meeting their needs of scalability, security and privacy. As an open source project based on Open Stack Swift, there is a large community that will be able to benefit from StackSync features. The project is also

open to third-party integrations like IaaS providers or software providers. This all will help organisations to face those key concerns that are actually impeding their adoption of the cloud.

Links and references

[1] www.stacksync.org.
[2] http://cloudspaces.eu.

Funding

CloudSpaces is a 3-year project funded under the European Commission's 7th European Framework, Software & Services, Cloud, DG CNECT (FP7 ICT Call 8).



SyncFree: Large-scale computation without synchronisation

Tyler Crain, Marc Shapiro, INRIA & LIP6

Large-scale on-line services built on top of massive cloud computing systems, including social networks and multiplayer games handle huge quantities of frequently changing shared data. A key requirement of these applications is to ensure data consistency, protection of confidential user information (e.g. protection of bank details) and appropriate permissions (e.g. only authorised sharing of information or content). Increasing scalability requirements are posing challenges to maintaining data consistency, which is becoming less straightforward in a centralised cloud. A key challenge lies in replicating data across several distributed datacentres, which requires new principled approaches to consistency.

From a European perspective, this is important because current evidence points to the multiplication of loosely coupled, widely distributed localised datacentres of all sizes. However, current solutions for ensuring data consistency in these systems require highly specialised, expert technology and investment, with application developments usually led by only a few large organisations.

SyncFree, a project funded under the European Commission's 7th Framework Programme, Software & Services, Cloud Computing, DG CNECT [1], is focused on addressing these challenges by driving new principled approaches to consistency.

Addressing key challenges for large-scale computation

The Internet is undergoing an incredible growth of interactive services involving millions of concurrent users, including scalable algorithms that provide weak or relaxed data sharing (e.g. MapReduce or Content Delivery Networks). However, many essential applications require robust sharing that maintains the consistency of shared and mutable data. Examples include massive multi-player online games, online mobile games, advertising platforms, collaborative social networks, and information networks (e.g. healthcare).

These applications have significant requirements in terms of the number of users, the amount of data, and geographical coverage. Maintaining strong consistency at this scale is becoming a major technological barrier for many online services because of issues like network delays, operational costs, and hardware failures.

The SyncFree project will develop scalable solutions to these problems, including libraries of open-source data structures available to developers of these services. SyncFree will address these challenges by using a recent, principled approach to enabling robust sharing, called Conflict-Free Replicated Data Types (CRDTs). CRDTs avoid the complexities of ad-hoc approaches, while maintaining the scalability advantage.

Here comes the insight. By following a few simple mathematical principle, for example commutativity, distributed updates can occur without synchronisation, while still ensuring a level of data consistency that enables the development of powerful applications. What's more, CRDTs ease development by encapsulating the replication and concurrency properties of common shared objects, such as sets, maps, sequences, or graphs.

Preliminary, small-scale experiments show that CRDTs have many advantages, such as locality of data, low latency of updates, and full-time availability, and that they require less computation and network resources. Maintaining consistency for real-world applications with millions of concurrent updates is where SyncFree comes into play.

Firstly, SyncFree will document the requirements of these types of applications, both in natural and in mathematical language, thus investigating their theoretical and practical scalability limits. A set of core CRDT algorithms will then be designed for these applications, studying trade-offs between scalability, consistency, and security, while examining the computational, network, and storage costs. Finally, the project will explore how to provide additional guarantees, such as transactional updates and bounded storage, which are required by certain applications, without impacting negatively on the advantages of CRDTs.

Who stands to benefit and how

The SyncFree project will advance both the theory and practice of large-scale application architectures, and especially of CRDTs and related mechanisms. As SyncFree partners from enterprise already have large user bases and feel the need for increased scalability in their applications, the project will include an extreme-scale crowd-sourced experiment, pushing the scalability needs of real world applications. An open-source library of CRDTs, to be used in future scalable distributed applications, will be made available, leaving a lasting and beneficial impact far beyond the end of the project.

Using these open source libraries, organisations will be able to create highly scalable programs more easily, thus meeting strict consistency requirements in today's highly connected services while improving user experience through low latency and fault tolerance. These advantages will help extend the reach of the cloud into mainstream connected applications and services.

Links and References

[1] https://syncfree.lip6.fr/.

Funding

SyncFree is funded under the European Commission's 7th European Framework, Software & Services, Cloud, DG CNECT (FP7 ICT Call 10). This 3-year project started in October 2013.

Standards Groups

DMTF distributed management task force, inc. IEEE OASIS OpenGridForum SNIA Europe

Media Partners

isgtw international science grid this week SC Online News