Zurich University of Applied Sciences

ZHAW School of Engineering

Institute of Applied Information Technology,
Service Engineering Research Area

Service Prototyping Lab Report - 2017 (Y2)

Authors:
Serhii Dorodko, Piyush Harsh, Manuel Perez Belmonte, Manuel Ramírez López, Oleksii Serhiienko, Martin Skoviera, Josef Spillner

September 6, 2017

Dear Reader,

This second lab report informs about another energetic year at the Service Prototyping Lab (SPLab) at Zurich University of Applied Sciences. A major focus has been on improving the scientific workflow with reference datasets, reproducible/repeatable/recomputable results, early preprints and transparent work communication through open notebook science. With such means, we believe to have reached a new level of research quality throughout the entire university in this regard and invite our fellow colleagues to join the ongoing discussion of how we as computer science community should work and publish in the future.
The bread and butter of lab activities remain technical research and innovation projects around the topic of delivering cloud applications. With three consolidated research initiatives on cloud-native applications, service tooling and cloud accounting and billing, the lab is now well-positioned to push further in this direction and discover new knowledge about how such applications should work and behave.
Apart from the research work, the lab has increased its commitment to education and is now responsible for filling one and a half elective modules with demanding and qualifying content. By educating the next generation of engineers and by working directly with local software industry, the Service Prototyping Lab contributes to the advancement of ICT also on a regional level with tangible transfers to companies. Our prototyping methodology produces new services and surrounding ecosystems rapidly to allow for analysis, evaluation and handover to production engineering.

Enjoy the lab report and find out more about our lab, its scientific publication, its curated data sets and its open source software products at the websites `http://blog.zhaw.ch/icclab`, `https://osf.io/khxk8/` and `https://github.com/serviceprototypinglab`, respectively.

Doz. Dr. habil. Josef Spillner
Senior Lecturer
Head, Service Prototyping Lab

The Service Prototyping Lab

One of the fastest-growing trends in terms of complex software-service systems is the use of cloud computing platforms to build and run distributed applications. Clouds are known to be hard to get into due to evolving APIs, even harder to get out of due to technology and vendor lock-in effects, and not without risks during their use. Still, cloud computing offers many compelling advantages to many application classes, including scalable web applications, industry applications such as smart connected things and robots, and mobile backends. The advantages encompass highly elastic scalability, flexible on-demand provisioning and wiring, fine-grained usage tracking and billing, and the outsourcing of critical tasks such as backups and updates. An additional advantage is the consequent use of service orientation to foster re-use of software services.

The *Service Prototyping Lab* concentrates research on overcoming the challenges to bring applications to the cloud with confidence in their predictable quality. The notion of prototyping incorporates the desirable properties *fast*, *low barrier* as well as *high quality* in additive combinations. The results are thus of benefit to companies in need to try out new technologies without high upfront investments in both training and software development, thus aligning with a truly *applied research* perspective. Through the Service Prototyping Lab, tools, guidelines, and system modifications will be propagated to streamline the process of onboarding applications into cloud environments.

Due to the wide range of topics around cloud applications and services, the lab is structured into three research initiatives around the topics of pervasive services and service-based applications: cloud-native applications, service tooling, and cloud accounting and billing.

The initiatives are long-living under the assumption of being led by a permanent researcher in the lab. Initiatives are fueled by funded projects, often running for shorter periods of time, as well as additional research activities. The initiatives will be presented on the subsequent pages. Research results are transferred to companies and into education alike which is a crucial element given the likely employment of students at local companies which are in turn candidates for knowledge and tooling transfers. The lab thus carries out research on how to bring applications into the cloud, considering the interfaces, tools, layers, processes and essential services.

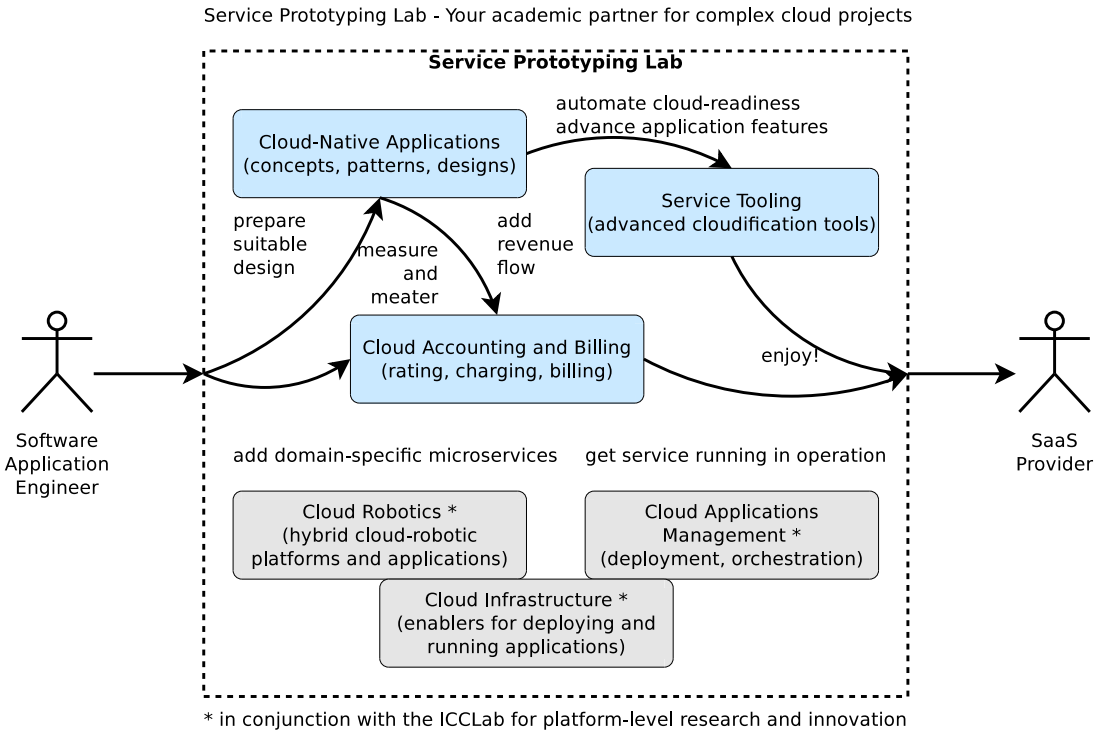

Interlinked research initiatives within the SPLab cover the entire cloud computing stack

Researcher Spotlight: Manuel Ramírez López

Manuel was born in El Burgo, a nice village in Málaga, Spain. He studied a bachelor in Computer Sciences and a bachelor and MSc in Mathematics, both in the university of Málaga.
He started to work in ARPA solutions, a company from Málaga where he worked inside Augmented reality projects. In May 2015, Manuel arrived to Switzerland with an IAESTE internship in the university FHNW and he was working for one year in the research projects FLARECAST and HELIO.
In June 2016, Manuel finished his studies and he joined to the SPLab as a scientific assistant where he solves software design problems in the cloud-native applications research initiative with a special interest in elasticity and stateful microservices. Within the initiative, Manuel initially started working in the ARKIS project where he designed and implemented a cloud-native microservices application which offers several multi-tenancy options or a testbed to compare cloud-native database. Currently, Manuel is working in the MOSAIC project to ease the onboarding of cloud applications in container platforms.

Related scientific publications

G. Toffetti, S. Brunner, M. Blöchlinger, J. Spillner, T. M. Bohnert: "Self-managing cloud-native applications: design, implementation, and experience", Future Generation Computer Systems (FGCS) 72 pp. 165–179, July 2017. DOI: 10.1016/j.future.2016.09.002

J. Spillner, G. Toffetti, M. Ramírez López: "Cloud-Native Databases: An Application Perspective", 3rd International Workshop on Cloud Adoption and Migration (CloudWays) @ ESOCC, Oslo, Norway, September 2017. To appear.

Research Initiative: Cloud-Native Applications

CLOUD research is shifting from the infrastructure and platform levels, which are largely commoditised by now, to the application level. A cloud-native application (CNA) is a distributed application that runs on a cloud infrastructure (irrespective of infrastructure or platform level) and is in its core scalable and resilient as well as adapted to its dynamic and volatile environment. These core requirements are derived from the essential characteristics that every cloud infrastructure must by definition possess, and from user expectations. It is of course possible to run an application in the cloud that doesn't meet all those criteria. In that case it would be described as a cloud-aware or cloud-ready application instead of a cloud-native application. Through a carefully cloud-native application design based on composed stateful and stateless microservices, the hosting characteristics can be exploited so that scalability and elasticity do not translate into significantly higher cost.
This research initiative offers design proposals, implementation guidance and evaluation testbeds for applied research with practical relevance. Furthermore, it monitors the fast-moving market for cloud application technologies including microservices, containers, composition and orchestration languages, and pricing trends for a holistic techno-financial assessment of cloud migrations.

Cloud-Native Applications research initiative

Related software

- `CNDBbench`. A benchmark for cloud-native database systems and services.
- `CNDBresults`. Reproducible experimental results when using CNDBbench.
- `ARKIS Microservices`. Cloud-native document management.
- `Scaling Containers`. Sample application with scalable stateless and stateful microservice.
- `Scalability Experiments`. Empirically finding boundaries of elastic scaling of containers.
- `Kube App Sizer`. Rightsizing Kubernetes applications.

Related blog posts

- "Cloud-Native Microservices Reference Architecture", July 2017
- "Rightsizing Kubernetes applications", June 2017
- "Cloud-Native Document Management", January 2017
- "Container management with Kubernetes: Practical Example", November 2016
- "Container Management with Vamp: Practical Example", November 2016
- "Benchmarking cloud-native database systems", September 2016

Related talks

- M. Ramírez López: "Predictable elasticity of Docker applications", 14th Docker Switzerland User Group Meetup, May 2017, Wallisellen, Switzerland

/ Cloud-Native Applications

Within the research initiative, research and innovation projects of relevance to the local industry are being worked on. In the recent months, the initiative has brought progress to two ICT companies from the regions of St. Gallen and Zurich in Switzerland.

Project Spotlight: ARKIS

Funded by the Swiss Commission for Technology and Innovation (CTI) and executed with Kendox AG, ARKIS aims at going beyond the current trend of "just" moving data and document management into the cloud by defining a cloud-native architecture for managing documents reliably and with scale in an ecosystem of third-party services around a document management system. Business-critical aspects such as fully compliant and auditable document access procedures are made service-oriented by rating and billing them on a per-use basis. Cloud features such as differentiated storage and surge pricing are analysed for their suitability in this particular domain.

Project Spotlight: MOSAIC

The MOSAIC project focuses on providing a platform for delivering any kind of application as a service, with a focus on container-based applications. It features an integrated incident management system as well as a container-optimized storage system. The platform will be able to deploy hybrid applications split into multiple locations, optimizing resiliency and cost in the process, as well as support continuous integration and deployment of each service.

Project MOSAIC aims to deliver a platform to deploy and manage distributed, container-based applications. None of the currently available Platform-as-a-Service frameworks provide the same benefits to application developers: MOSAIC delivers a vendor-independent, Platform-as-a-Service framework independent, software suite which can orchestrate applications on multiple providers, automatically monitor them during runtime, automatically detect and resolve runtime incidents, all based on a custom storage backend optimized specifically for container-based cloud-native applications. MOSAIC is funded by the Swiss Commission for Technology and Innovation (CTI) and executed with VSHN AG.

/ Cloud-Native Applications

Our research distinguishes between between legacy, cloud-enabled, cloud-aware and fully cloud-native applications. Legacy refers to applications which run only on premise or through manual installation procedures. Cloud enablement refers to the ability to deploy to private or public cloud environments in the form of virtual machines, containers, application packages, hosted functions or unikernels. Cloud awareness adds the ability to bind to platform services in the cloud environment through brokers or other standardised means, for instance to ensure database or e-mail sending functionality. Once applications are deployed, they may be re-deployed or migrated to another provider offering the same environment (homogeneous migration). Through transformation, other environments can also be targeted (heterogeneous migration). Finally, applications may become self-aware and migrate themselves to the most economic cloud provider (autonomic migration). Hence, we distinguish: 4 cloud application maturity levels, 5 representation technologies, and 3 levels of autonomy.
A further means of classification is the ability of applications to self-manage versus being managed by the hosting platform. Various trade-offs are involved, including different choices concerning multi-tenancy isolation.

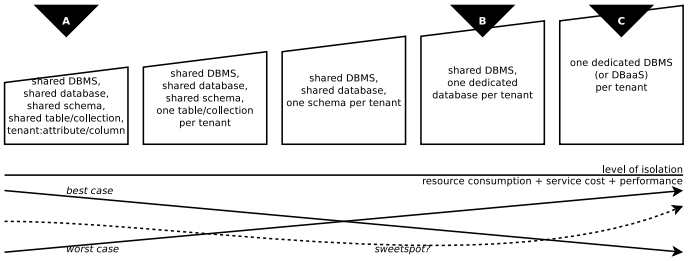

Five options for multi-tenancy in cloud-native applications

As application engineers are faced with such a high variety of options on their way into the cloud, the informed choice using quantified data from systematic tests becomes essential. Of special concern are stateful services whose practical implementations nowadays often do not adhere to the desired characteristics. This is why the Cloud-Native Applications research initiative has proposed, designed, implemented and evaluated a testbed for comparing self-managed and provider-managed database services.

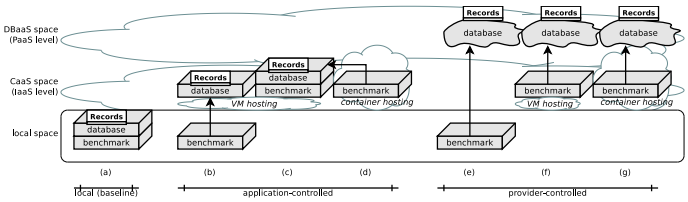

Systematic testing of cloud-native application characteristics

Using the testbed, generic reference implementations of composite microservice applications and specifically engineered testing tools, we have been able to explore the solution space for several legacy business applications for CRM and DMS.

Researcher Spotlight: Josef Spillner

Josef Spillner is affiliated with Zurich University of Applied Sciences as senior lecturer and head of the Service Prototyping Lab in conjunction with the InIT Cloud Computing Lab. Before founding the lab, he conducted research and led activities as post-doc at TUD, SAP, NTUU, UFCG and UniBZ. He wrote a doctoral dissertation about metaquality of services (2010) and a habilitation treatise about stealth computing in multi-cloud environments (2015) and published more than 50 papers on related matters. Josef Spillner is active in several cloud communities and initiated the Open Source Service Platform Research Initiative to promote re-usable software for scientific work.

Related scientific publications

J. Spillner, C. Mateos, D. A. Monge: "FaaSter, Better, Cheaper: The Prospect of Serverless Scientific Computing and HPC", 4th Latin American Conference on High Performance Computing (CARLA), Buenos Aires, Argentina, September 2017. To appear.

Related software

`MC-SIM/MC-EMU.` Multi-Cloud Simulation + Emulation framework for targeted failures of compute, storage and networking resouces.
`Transducer.` Service Interface Transducer for Rapid Prototyping.
`Whatcloud.` Identification of cloud provider by network location.
`Javacode.` Diverse Java projects with different build and unit test configurations.
`Lambda Control Plane.` Lambda Control Plane applications: Lambackup & LaMa.
`Podilizer Experiments.` Quantified knowledge about what Podilizer does.
`Podilizer.` Java to Lambda.
`Termite.` Lambdafication: Java to Lambda with annotations.
`Snafu.` The Swiss Army Knife of Serverless Computing.
`FaaS Experiments.` Testing AWS Lambda and other FaaS providers.

Research Initiative: Service Tooling

WORKING with remote services requires appropriate and decent tooling. A service idea may take just five seconds ("I want to offer a robust note-taking service"), but its realisation may take much longer ("Which programming language and model?", "How to describe the service?", "Where do I find a fitting file service to store the notes on unless I want to take care of backups by myself?", "Where do I publish my service so that it runs and generates income?"). Therefore, modelling, engineering and integration tools are primarily needed. These tools work in combination with a certain service environment, or ecosystem, consisting of more tools, dependency services, and service platforms which bring services to life.
The Service Tooling research initiative of the Service Prototyping Lab intends to identify tools and platform services which are straightforward to deploy, easy to use and generic enough to be re-usable in many service scenarios.
For this purpose, the initiative follows a triple structure with three topics of increasing industrial and societal interest: Function-as-a-Service (FaaS), Stealth Computing, Cloud Ecosystems.
In the FaaS part of the initiative, tools to bring legacy code into FaaS environments as well as tools to advance the environments themselves are investigated. There are software decomposition tools for Python (Lambada) and for Java (Podilizer, Termite). Furthermore, there is a flexible client/server tool to migrate, execute, test and deploy functions written in several languages (Snafu).

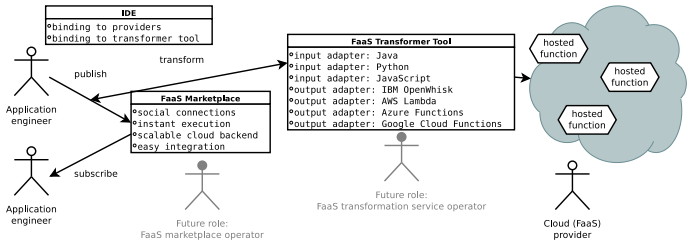

Envisioned FaaS ecosystem

The cloud ecosystems part of the initiative explores marketplaces, brokers, dashboards, cloud migration tools, API generators, aggregators and other enablers of thriving ecosystems with service producers and consumers. The research focuses on prototyping techniques with description/implementation roundtripping, a library of utility services which aid in establishing ecosystems, and improved client-side tools such as CLI helpers.

Related scientific preprints

- J. Spillner: "Transformation of Python Applications into Function-as-a-Service Deployments", arχiv:1705.08169
- J. Spillner: "Snafu: Function-as-a-Service (FaaS) Runtime Design and Implementation", arχiv:1703.07562
- J. Spillner, S. Dorodko: "Java Code Analysis and Transformation into AWS Lambda Functions", arχiv:1702.05510
- J. Spillner: "Exploiting the Cloud Control Plane for Fun and Profit", arχiv:1701.05945

Related blog posts

- "Termite. A Java library for the selective "Lambdafication" of applications", June 2017
- "Research Directions for FaaS", April 2017
- "Running Google Cloud Functions in OpenShift", April 2017
- "Snafu – The Swiss Army Knife of Serverless Computing", March 2017
- "Transducing service descriptions into SaaS prototypes", January 2017
- "Introducing Podilizer: Automated Java code translator for AWS Lambda", December 2016
- "Rapid API generation with Ramses", December 2016
- "Programmatic identification of cloud providers", December 2016
- "FaaS: Function hosting services and their technical characteristics", October 2016

Researcher Spotlight: Serhii Dorodko

Serhii is a final year student in Computer engineering at National Technical University of Ukraine "Kyiv Polytechnic Institute". He joined the SPLab through the IAESTE internship program and started to work on the 1st of September 2016. Serhii aims to get new experience and knowledge, improve programming and communicative skills, meet new friends and learn languages. This is his first employment and he is very inspired by potential to learn new skills both scientific and pertaining to software development. He is working in the Service Tooling initiative team. Due to this initiative he conducts research in advanced cloud technologies such as FaaS which is new and getting very popular.

/ Service Tooling

Finally, in the stealth computing part of the initiative, there are a number of architectures depending on the use case and the lifecycle phase of a service. The following diagram represents a typical multi-cloud service integration point with stealth properties. Software applications and services benefit from spreading their data and functions across providers in a tightly controlled, re-usable layer with standard interfaces such as files (e.g. POSIX) and data (e.g. SQL). Users are more willing to adopt cloud environments when explicit user control is made possible by stealth computing.

Service Tooling research initiative

The focus in the recent months was clearly on FaaS due to the emerging industrial trend towards serverless computing. Among the work conducted was an initial survey about the ecosystem landscape which includes programming conventions, public FaaS runtime providers, open source stacks for private runtimes, development abstraction frameworks, developer tools to deploy and debug, and integration of FaaS stacks into existing cloud stacks.
Our major contributions in this space have been code transformation tools to bring legacy code into serverless environments, and the Swiss Army Knife of Serverless Computing, a versatile tool and prototypical runtime to manage and run functions locally and across various major cloud providers. We are now in the process of transferring the results to local SMEs.

Related talks

J. Spillner: Serverless Applications: Tools, Languages, Providers and (Research) Challenges, Serverless Zürich, June 2017, Zurich

J. Spillner: Function-as-a-Service: A Pythonic perspective on Serverless Computing, PyParis, June 2017, Paris, France

J. Spillner: More on FaaS: The Swiss Army Knife of Serverless Computing, Future Cloud Applications #2, April 2017, Winterthur, Switzerland

J. Spillner: Containerising Functions using Docker and OpenShift, Microservices Zürich, April 2017, Zurich, Switzerland

/ Service Tooling

To enable global research on FaaS and around serverless technologies, we have evaluated the fitness of this programming and deployment model for scientific work, aiming at scientific computing environments and high-performance computing. Using Snafu and carefully engineered portable hosted functions, we were able to compare cloud-provided FaaS with self-hosted FaaS concerning the use of compute, storage and network resources. Surprisingly, most providers only offer a vague instance configuration through the assignment of available memory to function instances. This contrasts the plethora of instance types for hosted virtual machines. We therefore expect a much more refined serverless offering accounting for such differences in the future.

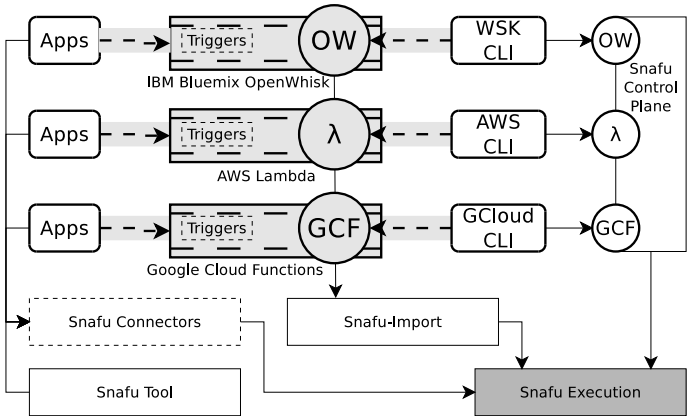

Integration of Snafu into the major cloud provider ecosystem of serverless computing

Several performance results from four different scientific computing domains could be obtained which suggest that future distributed application development will become more straightforward with function-oriented programming constructs such as function futures.

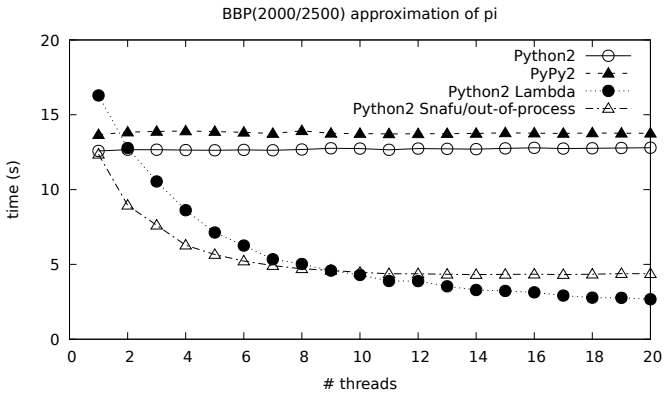

Seemingly serverless approximation of π

Researcher Spotlight: Piyush Harsh

Piyush Harsh is a researcher in ZHAW Service Prototyping Lab. His research interest includes distributed self-* capable systems, nature inspired computing, authentication protocols, cloud computing, to name a few. Before joining SPLab, Piyush was the development coordinator of WP5 in FP7 project Contrail and was instrumental in the development of cloud-middleware tool Virtual Execution Platform.
Piyush received his bachelors in Computer Science from Indian Institute of Technology and then received his MS and PhD degrees in Computer Engineering from Department of Computer and Information Science and Engineering at University of Florida. In the past Piyush worked as research engineer at INRIA Rennes Bretagne-Atlantique research center.
At SPLab, Piyush is involved in FP7 IPs Mobile Cloud Networking and T-Nova projects.

Related scientific publications

- M. Skoviera, P. Harsh, O. Serhiienko, M. Perez Belmonte, T. B. Bohnert: "Monetization of Infrastructures and Services", European Conference on Networks and Communications (EuCNC), Oulu, Finland, June 2017

Related software

- `Cyclops.` Comprehensive rating-charging and billing solution for cloud services.

Related blog posts

- "Openstack checkpointing is simplified", January 2017
- "Cyclops 2.0 Dashboard is Out!", November 2016
- "Integration of Openstack OVA importing tool to Horizon", October 2016
- "A new tool to import OVA Applications to Openstack", September 2016

Research Initiative: Cloud Accounting & Billing

FINANCIAL accounting is a very critical process in the monetization process of any service. In the telecommunication world, these processes have long been documented, used, and standardized. Cloud computing being a relatively new paradigm, is still undergoing a transition phase. Many new services are being defined and there is still a huge untapped potential to be exploited.
Rating, Charging, and Billing (RCB) are key activities that allows a service provider to fix monetary values for the resources and services it offers, and allows it to bill the customers consuming the services offered. The Cloud Accounting & Billing research initiative transforms this potential into concrete technical solutions for small and medium enterprises. The challenges we are trying to address under this initiative are multi-dimensional. Is it possible to come up with a general enough RCB model that can address the needs of multiple cloud services – IaaS, PaaS, SaaS, and many more that would be defined in the future?

Project Spotlight: Scale-UP

swissuniversities

In SCALE-UP, a Distributed Computing Platform (DCP) in the cloud will be developed that will allow researchers to provision and use their favourite data processing framework (among a multitude including Hadoop, Spark, Storm, etc.) on-demand over popular cloud frameworks. For researchers who may not know the most suitable framework for their data processing needs, the DCP will provide a matching algorithm that will guide the researcher to the most suitable one interactively. SPLab will bring in their advance rating-charging-billing framework Cyclops into SCALE-UP. Cyclops will be integrated into existing cloud infrastructure offered by SWITCH. The goal is to enable SWITCH and partner institutions (Swiss universities, research labs) implement most-suitable pricing and billing models for any tools and services they offer to the community. This will significantly advance the self-sustainability of various products and services being developed in the project. Towards this target, a marketplace is envisioned and Cyclops framework will be customized to support self-configuration of billing models for products and services. SCALE-UP is funded under a CUS programme by Swissuniversities.

Related talks

P. Harsh: Cyclops 3.0 – Hierarchical billing made simple for future cloud applications, Future Cloud Applications #3, July 2017, Zurich, Switzerland

O. Serhiienko: DMTF OVA on-boarding tool integration in OpenStack, 14th Swiss OpenStack User Group Meetup, November 30, Zurich, Switzerland

Researcher Spotlight: Martin Skoviera

Martin Skoviera is a researcher at SPLab, InIT ZHAW, focusing on the Rating-Charging-Billing initiative and its OpenSource solution – Cyclops.

Martin received his Bachelor's in Computer Science and Master's degrees in Management and Information Technologies from Brno University of Technology. Martin also has an MBA specialising in International Business and Trade Management from the Ligs University, Prague. He brings his technical as well as business expertise to the RCB team.

In the past Martin was actively involved in the banking sector, developing mobile solutions, and working with both ITIL and COBIT management frameworks. Martin's first introduction in ZHAW was as an exchange student, where his work centred around Business Engineering and Production.

Martin is now working on the SafeSwissCloud and ScaleUP projects, as well as contributing in the lab's EU projects – TNova and MCN.

/ Cloud Accounting & Billing

Any successful monetisation of application functionality requires a well-connected monetisation process. Functionality is first metered, then processed through mediation and accounting rules, and finally matched against a pricing model to generate proper charges. Real-time or aggregated bills will be generated based on these charges. As a last step, the process includes financial clearing to ensure the used functionality has been paid for.

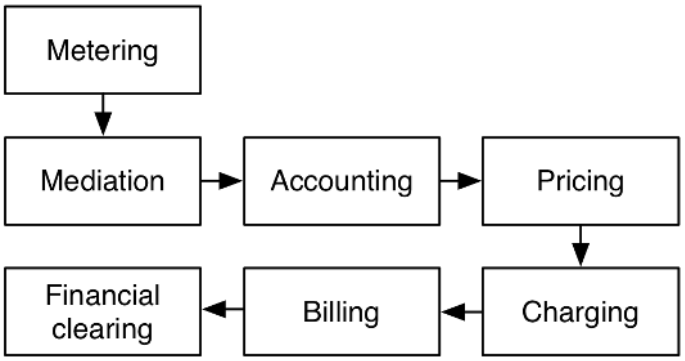

Monetisation process for rating, charging and billing of cloud services

One of the early results of this initiative has been Cyclops. It is a comprehensive dynamic rating-charging and billing solution for cloud services designed for IaaS, as well as PaaS and SaaS. In Cyclops, a UDR micro-service connects to cloud installation such as OpenStack, CloudStack or applications and autonomously collects the usage data. Cyclops enables the majority of the indicated monetisation process.

Harmonised data is being saved into a time series database and exposed through REST APIs. Subsequently, dynamic rates are generated through a rule engine by considering the various policies, environmental factors and pricing strategies. A charge for the usage made is calculated by using the harmonised data from UDR micro-service. Finally, various plans, discounts and coupons are considered through a billing engine to calculate the price to be paid by the user. Once the invoice is generated, the service connects to the payment gateway for payment fulfilment. Data from the UDR, RC and Billing micro-services is presented to the end user in the form of charts, statistics and bills. Administrators manage user accounts, select what meters to use and configure the rating details.

Researcher Spotlight: Manu Perez Belmonte

Manu is a Computer Engineering student from Mataró, Barcelona. He is currently in his final year of studies at the Technocampus, Barcelona. He joined the SPLab through the IAESTE internship program and is working as an intern from the 1st of July.
He is very interested in learning new things, news technologies, meeting friends, sports and learning new languages.
This is the first study-related work experience for Manu and he will be working in the Rating Charging Billing initiative team.

Researcher Spotlight: Oleksii Serhiienko

Oleksii Serhiienko is a part-time master student and researcher at the SPLab working on the on the Rating-Charging-Billing initiative and its OpenSource solution –Cyclops.
Oleksii has been graduated at the Kiev Polytechnic University majoring computer engineering. Before he had already been as a part of the SPLab community. In 2014, he resided in the laboratory like exchange student by IAESTE program. During that year, he was working on OpenStack in particular at the Ceilometer projects. His code has been added to the "Icehouse" OpenStack release. When he came back after internship to Ukraine he continued working with OpenStack technologies and python programming.
Currently, Oleksii is working on the SafeSwiss cloud project in particular on the Prediction engine development.

/ Cloud Accounting & Billing

The Cyclops architecture consists of several composed microservices which allow for a reliable use in production environments. In tests with up to 16 virtual machines, the monetisation workflow has been able to ingest data at an average volume of almost 12x the single machine volume, suggesting a rather modest loss of linear scalability in line with Amdahl's law. On the output side, it has been taken no more than 12 seconds to produce bills with 50,000 entries for a customer across 50 event streams and 800 sessions.

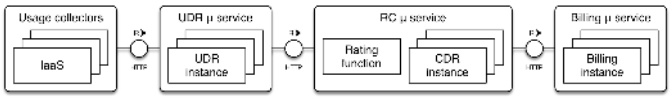

Microservice architecture for rating, charging and billing of cloud services

Future work focuses on flexible cost models which enable the prediction of prices for both application and hosting providers, leading to optimal matches and less frustration due to overcharging. From a research perspective, forecasting, optimised handling of time series and automated mapping of provider-informed pricing conditions to instances of generic models remain open challenges. Pricing strategies such as time-based, priority-based, Paris metro, edge pricing, proportional fairness and flat rate furthermore present a wide variety whose unification into models which can still be managed and validated remains an open question.

Education

TEACHING the basic elements of cloud computing found in several textbooks is not enough. We continuously update our lecture and lab materials based on findings in our research. Students benefit from this approach by getting some of the strongest education in the field of service and cloud computing.

Guest lectures in Paraguay

In August 2017, Josef Spillner gave two guest lectures on cyber-physical and cloud-native applications at the Itaipu Technology Park (PTI) and at the National University in Asunción, respectively. Furthermore, an intensive practical course over three half days was given at PTI using a flipchart, live demos and hands-on sessions. The remaining time was used to discuss ideas for cloud-native and resilient software services designs. All lectures, courses and discussions were conducted in Spanish. Partial funding was provided by PTI. Following this invitation by Fabio López Pires, we are currently planning to host two guest researchers from PTI at the Service Prototyping Lab next year.

Inside Parque Tecnológico Itaipu

Blog post: "Paraguayan Perspective on Cloud Applications", August 2017

Internet Service Prototyping (bachelor, elective module)

The elective module Internet Service Prototyping, offered in English in the 5th semester for both Computer Science and Business Engineering students, has been chosen by students again for the second iteration. The lecture conveys techniques on how to build cloud applications quickly without compromising quality. It introduces new tools like Vamp, the RAML workbench, Cyclops and distributed key-value stores to support the prototyping.
Blog post: "Reflections on Teaching Internet Service Prototyping", December 2016

Cloud Computing 2 (bachelor, elective module)

Our research is rapidly transferred to education. As part of this setting, the Service Prototyping Lab is responsible for a part of the Cloud Computing 2 lectures, labs and exams. Relevant topics include cloud-native patterns, service brokers and billing architectures as well as the interfacing with cloud platforms (PaaS).

Lab Life and Qualification

BEYOND the central research projects and endeavours clustered around our initiatives, the Service Prototyping Lab engages in qualifying the next generation of researchers and technologists.

Research and publishing approach

⚙ Blog post: "When the Open Access world is not enough", February 2017
⚙ Blog post: "Cloud Services: An Academic Perspective", December 2016
⚙ Blog post: "First Year of Service Prototyping Lab", September 2016

Doctoral theses

SPLab co-supervises two doctorate propositions in its areas of expertise. Stanley Ramalho Lima, from Brazil, spent half a year at Technische Universität Dresden, Germany, followed by joining the doctoral programme of University of Coimbra, Portugal, to investigate matters of resilience and complex faults in cloud computing environment. The research is supported by the Brazilian Ciências sem Fronteiras programme.
Ambrósio Vumo, from Mozambique, is currently spending half a year in Cologne, Germany, to learn German intensively. Subsequently, he will join Technische Universität Dresden as a sandwich doctoral student with Universidade Eduardo Mondlane, Mozambique, to revisit the state of cloud computing and networking in his home country.
⚙ Scientific publication: A. P. Vumo, J. Spillner, S. Köpsell: "Analysis of Mozambican Websites: How do they protect their users?", 16th International Information Security South Africa Conference (ISSA), Johannesburg, South Africa, August 2017

Apprenticeship

Janine Walther has concluded her activities in the SPLab as part of her apprenticeship on computer engineering for application development (Informatikerin Applikationsentwicklung). The project management SaaS developed by her was extended with a PDF-to-text extraction so that timesheets sent by e-mail can be imported into the system. This work has allowed her to pass the final defense.
⚙ Blog post: "Project Management SaaS", April 2017
⚙ Software: "PMSaaS: Project Management SaaS", April 2017


Left: Successful apprenticeship defense presentation; right: Demos at Open Cloud Day 2017

SPLab alumni

Janine Walther: After finishing her apprenticeship, she has moved into regional employment in the technology sector. Giovanni Toffetti Carughi and Tobias Lötscher: Both continue to work on the Cloud Robotics research initiative, now under the wraps of the ICCLab.

Events

EXCHANGING ideas with fellow researchers on an international scale is a central aspect of our work. For this purpose, we participate not only as attendees but also as organisers in events around the world.

Presence at events

SPLab participated in the following events during the reporting period with talks and technical demonstrations.

- Open Cloud Day, June 2017, Bern, Switzerland
- European Conference on Networks and Communications (EuCNC), June 2017, Oulu, Finland
- 1st PyParis conference, June 2017, Paris, France
- 3rd International Conference on Networked Systems (NetSys), March 2017, Göttingen, Germany
- 5th ROSCon, October 2016, Seoul, Korea

⚙ Blog post: "PyParis'17 conference report", June 2017
⚙ Blog post: "NetSys'17 conference report", March 2017

Local events

Furthermore, beyond participating in local technology meetups, the lab has successfully established its own hybrid meetup/evening seminar series called Future Cloud Applications (FCA). The first two events, FCA #1 and FCA #2, were organised in the ZHAW School of Engineering in Winterthur, while FCA #3 happened in ZHAW Toni-Areal in Zurich with talk contributions by Piyush Harsh (SPLab) and Evangelos Pournaras (ETH Zurich). The FCA series will be continued with further editions on demand to foster the collaboration between academia and industry.

Talk by Evangelos Pournaras on Engineering Democratization in Internet of Things Data Analytics

Community

ACADEMIC community services are important elements of self-administration. Through the operation of the cloud-announce mailing list and the active participation as committee members and reviewers, the Service Prototyping Lab is strongly involved in advancing cloud applications and services research.

Cloud-Announce mailing list

SPLab has been operating the cloud-announce mailing list with around 200 subscribers for almost one and a half years. Calls for papers, conference calls and research job offers are in its scope. The list is popular due to its moderate amount of e-mails without duplicates. The following figure compares submissions and approved posts. Stay informed by subscribing at `https://mailman.engineering.zhaw.ch/listinfo/cloud-announce`.

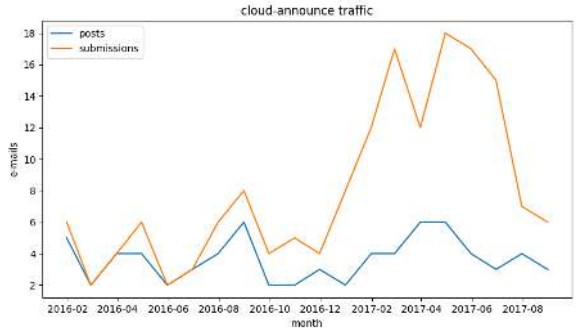

Cloud-announce statistics since February 2016

Blog post: "Cloud-Announce Statistics and Policy", January 2017

Young investigators

Within its second year, SPLab gained two very young investigators who have immediately started to explore the world and share their exciting findings with everybody around. Polina Spillner was born in November 2016 and Dhruv Harsh in June 2017.

Polina Spillner reading a book on distributed systems by Luntovskyy and Spillner, and Dhruv Harsh exploring containerisation technologies

Summary of Research Outputs

BEING an applied research laboratory implies a broad definition of long-term archivable research outputs. As we aim for high standards concerning research methods, ethics and dissemination, we engage in public preprints with public reviews, a positioning of selected software and data repositories as outputs as well as funding transparency.

Peer-reviewed publications (5)

M. Skoviera, P. Harsh, O. Serhiienko, M. Perez Belmonte, T. B. Bohnert: "Monetization of Infrastructures and Services", European Conference on Networks and Communications (EuCNC), Oulu, Finland, June 2017
G. Toffetti, S. Brunner, M. Blöchlinger, J. Spillner, T. M. Bohnert: "Self-managing cloud-native applications: design, implementation, and experience", Future Generation Computer Systems (FGCS) 72 pp. 165–179, July 2017. DOI: 10.1016/j.future.2016.09.002
A. P. Vumo, J. Spillner, S. Köpsell: "Analysis of Mozambican Websites: How do they protect their users?", 16th International Information Security South Africa Conference (ISSA), Johannesburg, South Africa, August 2017
J. Spillner, G. Toffetti, M. Ramírez López: "Cloud-Native Databases: An Application Perspective", 3rd International Workshop on Cloud Adoption and Migration (CloudWays) @ ESOCC, Oslo, Norway, September 2017. To appear.
J. Spillner, C. Mateos, D. A. Monge: "FaaSter, Better, Cheaper: The Prospect of Serverless Scientific Computing and HPC", 4th Latin American Conference on High Performance Computing (CARLA), Buenos Aires, Argentina, September 2017. To appear.

Preprints (4)

J. Spillner: "Transformation of Python Applications into Function-as-a-Service Deployments", arχiv:1705.08169
J. Spillner: "Snafu: Function-as-a-Service (FaaS) Runtime Design and Implementation", arχiv:1703.07562
J. Spillner, S. Dorodko: "Java Code Analysis and Transformation into AWS Lambda Functions", arχiv:1702.05510
J. Spillner: "Exploiting the Cloud Control Plane for Fun and Profit", arχiv:1701.05945

Data repositories and testbeds (4)

J. Spillner, S. Dorodko: "Lambdafication Repeatability", osf.io/c886p
J. Spillner: "FaaSter, Better, Cheaper", osf.io/8qt3j
J. Spillner, M. Ramírez López: "Quantifiable Scaling", osf.io/6gup8
J. Spillner, A. Vumo: "Mozambican Websites", osf.io/35sz8

Transparency declaration

All project and collaboration funding sources are indicated in this report. Piyush Harsh and Martin Skoviera are involved with Cyclops Labs GmbH, a spin-off from the Zurich University of Applied Sciences.