

ZURICH UNIVERSITY OF APPLIED SCIENCES

ZHAW SCHOOL OF ENGINEERING

INSTITUTE OF APPLIED INFORMATION TECHNOLOGY,  
SERVICE ENGINEERING RESEARCH AREA

---

## Service Prototyping Lab Report - 2018 (Y3)

---

*Authors:*

Mohammed Ali Mohammed AL-AMEEN, Serhii DORODKO, Piyush HARSH, Stefan JUNKER, Manuel PEREZ BELMONTE, Sebastiano PANICHELLA, Manuel RAMÍREZ LÓPEZ, Oleksii SERHIENKO, Josef SPILLNER

*Editor:*

Josef SPILLNER

August 23, 2018



**Dear Reader,**

It is my pleasure to report on the third year activity of the Service Prototyping Lab at Zurich University of Applied Sciences. A total of *nine people* contributed to our sustained impactful research during the last twelve months, of which seven continue into the fourth year, including now *two permanent senior researchers*.

One of the key characteristics of this year has been a strengthened international networking through incoming visiting researchers to our premises and guest speakers in our new colloquium series as well as outgoing activities, leading to first collaboration agreements. We also afforded more intensive travelling to meetings and conferences by lab staff based on convincing technical achievements. Along with our increased publishing activities, this sets us apart as preferred research centre for cloud applications and service development. Indeed, we have maintained our aim to stimulate community-wide discussions about modern forms of research, and have ourselves participated with registered datasets and preregistrations to be well-prepared to upcoming funding body requirements.

Furthermore, the conclusion of the third year signals the end of the planned three-year incubation period. All lab structures and qualified staff are in place to tackle the next challenges in fully exploiting the capabilities of cloud platforms for innovative global-scale distributed services. We are particularly proud of our new web home at <http://blog.zhaw.ch/splab>.

Bookmark this website and attend our events in Zurich in December - 11th IEEE/ACM UCC, 5th IEEE/ACM BDCAT, and the European Symposium on Serverless Computing and Applications (<http://essca2018.servicelaboratory.ch/>)!

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

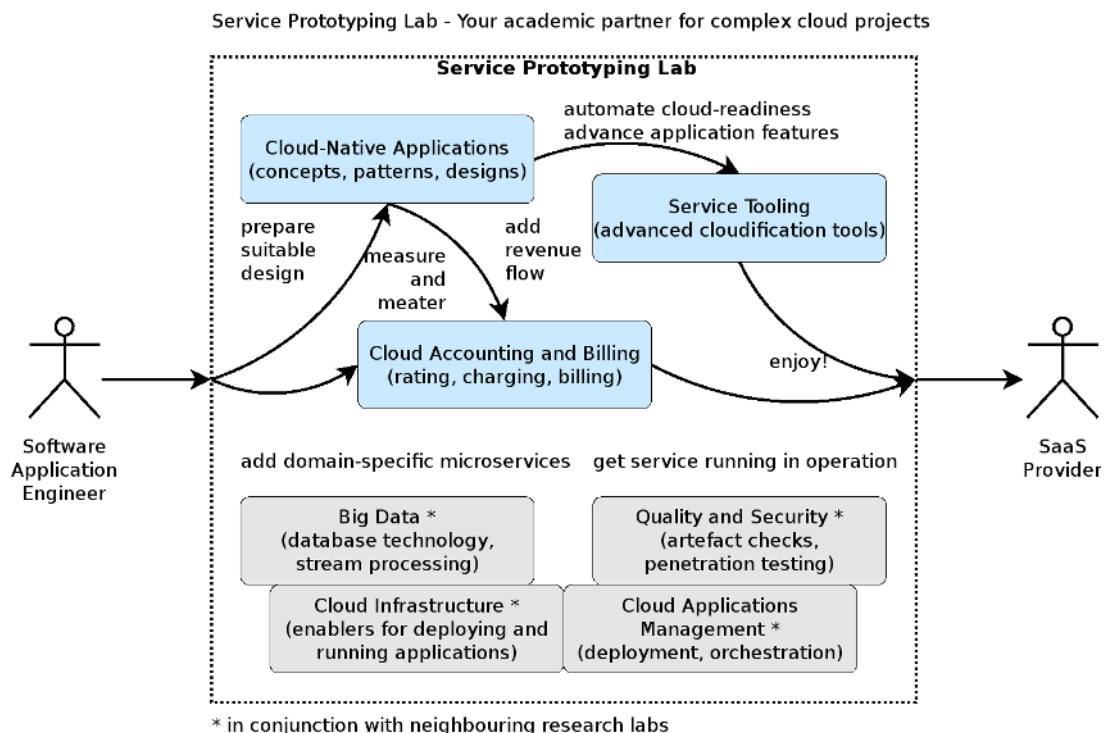
## The Service Prototyping Lab

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 development lifecycle*

## SPLab Colloquium Series

RENOWNED researchers from around the world are given the opportunity to present their work, views and opinions in the Service Prototyping Lab colloquiums. In the recent months, we have enjoyed and highly valued the presence of speakers from Poland, the US, Argentina (via Italy) and Germany. Maciej Malawski has explained the intrinsic advantages and disadvantages of serverless application architectures in the domain of scientific computing. His team at AGH University of Technology in Kraków, Poland, has conducted long-term monitoring of cloud providers to back up the claims. Further on the spectrum between scientific and high-performance computing, Alan Sill from Texas Tech University, USA, has shared his views on how data centre management and in fact application-controlled cloud resource management should be designed and implemented with recent standards and technologies. Back to cloud functions, Martin Garriga from Politecnico di Milano has outlined the use of these functions on mobile devices, on the edge and in the cloud at the same time. Finally, Christopher Hauser from Ulm University explained current cloud research. We are glad to have had these guests with us and welcome further speakers in our colloquium series.



*Impressions from the first colloquium with Maciej Malawski*



*Colloquium with Alan Sill*

- ⚙️ Blog post: “SPLab Colloquium on Serverless Scientific Computing“, February 2018
- ⚙️ Blog post: “SPLab Colloquium on Robust Modern API Design“, February 2018
- ⚙️ Blog post: “SPLab Colloquium on Serverless Continuum“, April 2018
- ⚙️ Blog post: “Colloquium on Cloud Computing“, August 2018

## Researcher Spotlight: Sebastiano Panichella

Sebastiano Panichella is a senior researcher who joined SPLab in August 2018 after his post-doc period at the University of Zurich. He is primarily responsible for leading the Cloud-Native Applications research initiative, including novel aspects such as continuous testing, integration and delivery of cloud applications as well as analysing software artefacts used in their construction. Among other tasks, he co-chairs the CNAX 2018 workshop. Sebastiano brings several years of EU and SNF project experience and a proven track record of high-quality publications in the area of software engineering into the lab.



## 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.

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. In his second year, Manuel has been working in the MOSAIC project to ease the onboarding of cloud applications in container platforms. He left SPLab in June 2018.



## Research Initiative: Cloud-Native Applications

CONTAINERISED applications are commonplace, the Cloud Native Computing Foundation releases an increasing number of industrially relevant software solutions, and DevOps processes are refined to actually work in cloud application development projects. Where is the research perspective in all of this? In our research initiative, we have focused on application migration in both senses of the word: Migrating (modernising) a legacy application through transformation and re-engineering, and migrating (moving) an application from one platform to another.



*Cloud-Native Applications research initiative*

### From legacy to cloud-native

In joint work with the Itaipu Technology Park, we have modernised the architecture and implementation of HENDU, a music royalty management application. The work was published at CLOSER 2018. We foresee the need to let a fully cloud-native application determine autonomously which functionality is self-managed and which one is entrusted to the underlying platform. For example, a hosted database service may be activated dynamically, for instance through the de-facto standard Open Service Broker API, with high reliability but also high cost, and therefore an application may want to run a self-hosted database in a container within the application context for some workloads.

## Related scientific publications

✿ M. Ramírez López, J. Spillner: “Towards Quantifiable Boundaries for Elastic Horizontal Scaling of Microservices“, 6th International Workshop on Clouds and (eScience) Applications Management (CloudAM) / 10th International Conference on Utility and Cloud Computing Companion (UCC), Austin, Texas, USA, December 2017. DOI: 10.1145/3147234.3148111

✿ J. Spillner, Y. Bogado, W. Benítez, F. López Pires: “Co-transformation to cloud-native applications: Development experiences and experimental evaluation“, 8th International Conference on Cloud Computing and Services Science (CLOSER), Funchal - Madeira, Portugal, March 2018. DOI: 10.21256/zhaw-3508

## Related software

- ✿ OS2OS. OpenShift application migration.
- ✿ Composeless. Hybrid cloud application composition.
- ✿ Volume2Volume. Migration of application data between OpenShift clusters.

## Related blog posts

✿ “Migrate OpenShift applications with os2os“, September 2017

## Related talks

✿ M. Ramírez López: “Migration of OpenShift Apps with OS2OS“, 2nd Cloud Native Computing Meetup, Zurich, Switzerland, September 2017


✿ M. Ramírez López: “Migrating data in a containerized application with volume2volume“, Future Cloud Applications Meetup #4, Zurich, Switzerland, April 2018

✿ J. Spillner: “Cloudware and Beyond: Engineering Methods and Tools“, 7th Open Cloud Day, Winterthur, Switzerland, May 2018

✿ J. Spillner: “Understanding Mixed-Technology Cloud Applications“, AGH University of Science and Technology, Kraków, Poland, June 2018

## / Cloud-Native Applications

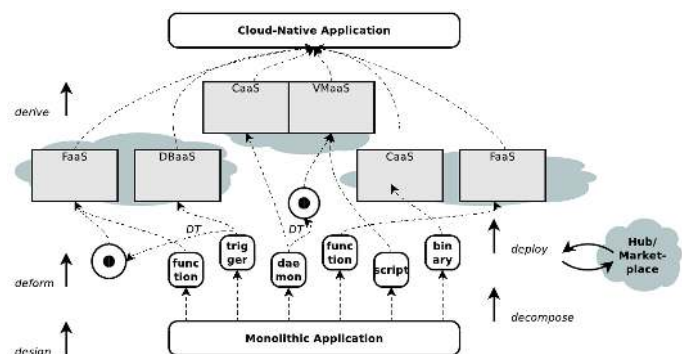
The dynamic self-management is shown in the following figure. Evidently, supporting such a design requires strong tooling support to avoid the need for hand-coding all possible combinations and checking their validity in any deployment configuration.



*Dynamically self-managed application*

## Mixing technologies

While containers are omnipresent, cloud functions are increasingly popular, and virtual machines are still around. There are currently no convincing concepts to combine these execution technologies coherently during development, composition and execution. We have created a proof-of-concept prototype, Composeless, based on Docker Compose to explore possible designs. Further, we have created a concept for decomposing applications into mixed-technology artefacts, as is shown in the figure below. The realisation is not trivial, but feasible considering our previous work on semi-automated decomposition of monolithic code into individual functions, controlled by annotations on the code level.



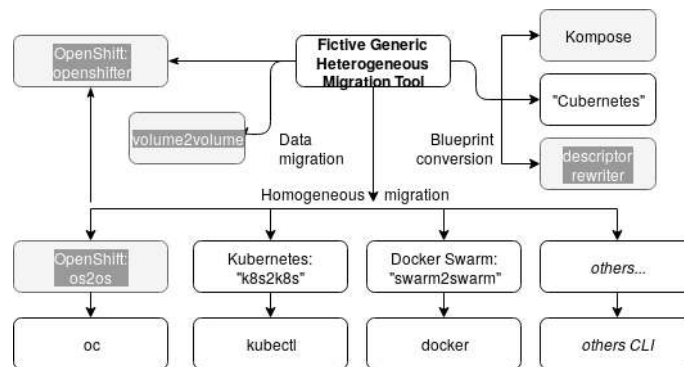
*Mixed-technology artefact decomposition*

With a cloud-native mindset, one should then look further and envision software development environments which suggest the appropriate runtime technology for each section of the code, while supporting the right level of granularity and the right amount of coupling. Certainly, with vendors like Eclipse going cloud-native, cloud computing and software engineering coming closer together will finally extend to industry, too.

## / Cloud-Native Applications

### From one platform to another

Moving (or, related, copying) composite microservice-based cloud applications requires moving the composition description(s), moving or redeploying the microservice implementation, and moving the data. We were concerned primarily about two challenges: The inhomogeneity between platforms, if not even full heterogeneity in some cases, as well as the correct transactional handling of data. Several prototypes were created to explore the problem specifically for OpenShift container platforms: OS2OS for the services, Volume2Volume for the data, and OpenShifter as next-generation conceptual sketch which can be run as a service within the source or target environment itself. The following figure shows a typical migration use case which can be realised through these tools, even extending to heterogeneous cases with additional transformation tools.



*Cloud application migration in heterogeneous environments*

### 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 Innosuisse and executed with VSHN AG.

## Researcher Spotlight: Josef Spillner

Josef Spillner is a senior lecturer and head of the Service Prototyping Lab at Zurich University of Applied Sciences in Switzerland. His research interests include cloud-native applications, service tooling and cloud accounting & billing. With his team, he works on challenging topics such as microservices, function-as-a-service and time series analysis. Before founding the lab, he conducted research at TUD, SAP, NTUU, UFCG and UniBZ and founded the Open Source Service Platform Research Initiative to promote re-usable software for scientific work. His work approach continues to promote international exchange, most recently at PTI and IIT, and modernisation of research and publication approaches. He published a doctoral dissertation about metaquality of services and a habilitation treatise about stealth computing in multi-cloud environments.



## Researcher Spotlight: Mohammed Ali Mohammed Al-Ameen

Mohammed Al-Ameen is a research intern on visit from the OpenUAE group at University of Sharjah, United Arab Emirates. His contributions at the Service Prototyping Lab thus far encompass assistance to conference organisation and refactoring quality assessment software. Among his research interests are serverless computing architectures and Internet of Things.



## Research Initiative: Service Tooling

IN recent months, industrial-strength tools to develop and manage services have emerged but are still lacking on many fronts. From our research perspective, we had particular interest in the quality dimension. What happens when several software artefacts, each with small quality deficiencies, are composed into a complex application? The result can easily be a totally unusable product. Therefore, we started to monitor digital artefact ecosystems and scan and analyse the artefacts for quality deficiencies.



*Service Tooling research initiative*

## Ecosystem monitoring

In academia, service ecosystems have been envisioned a long time ago, but only now are they really emerging in industrial practice. Hubs with thousands or millions of software artefacts are becoming a one-stop shop for a new generation of developers who are already familiar with glueing together snippets, for the better or the worse. We have conducted a regular monitoring of two repositories over three months each: AWS Serverless Application Repository which contains deployable cloud functions and KubeApps Hub which contains deployable Helm charts as enclosures of Kubernetes deployment descriptors. The monitoring continues, but the intermediate results are already interesting per se based on the data which we compiled and published in the Open Science Framework.



## Researcher Spotlight: Stefan Junker

Stefan Junker has been a research assistant in the Service Prototyping Lab at the Zurich University of Applied Sciences. After finishing his Information Technology specific Abitur he started gathering industrial experience during a three-year apprenticeship as an IT specialist. Afterwards he began his academic career with a B. Sc. and M. Sc. both in Computer Sciences at the HTWG Konstanz. His B. Sc. Thesis was on Car to Car communication services, where he worked on service discovery and messaging in meshed WiFi networks. In his M. Sc. Thesis, he evaluated the effect on memory-safety when using a memory-safe programming language (Rust) to implement the multi-tasking features within an Operating System. Parallel to his studies, he worked as an IT laboratory engineer at HTWG Konstanz, as a self-employed IT specialist supporting customers with cloud-based virtual desktop environments, and contracted for CoreOS Inc. to develop cutting-edge Linux-container technologies. Towards the end of his studies, he had the opportunity to hold lectures on Operating Systems and System Software to Bachelor students at HTWG Konstanz.

After finishing all the above, he joined the Service Prototyping Lab due to his interest in API driven Cloud Services, where he has started to conduct research in intermediary services in multi-cloud environments. He left SPLab in July 2018.

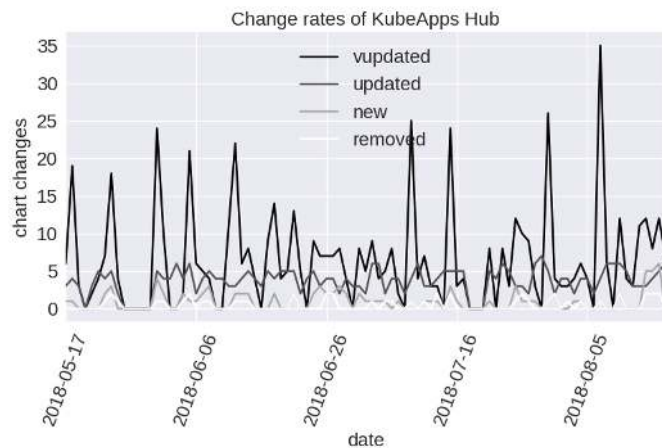


### Related scientific preprints

🔧 P. Leitner, E. Wittern, J. Spillner, W. Hummer: “A mixed-method empirical study of Function-as-a-Service software development in industrial practice“, PeerJ Preprints 6:e27005v1, June 2018.

## / Service Tooling

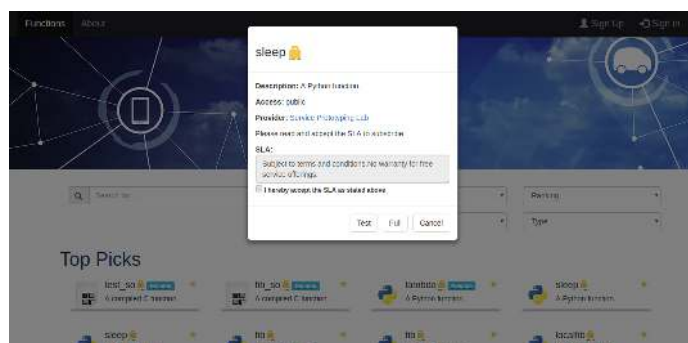
The following figure shows, for example, the frequency with which minor updates as well as semantically different version updates are performed on the about 200 Helm charts curated in KubeApps Hub. More importantly, however, are metrics related to quality: How many maintainers are involved in a chart, how many variable and duplicate values are found in the templates, is a source code management URL specified for opening issues automatically or an e-mail address for contacting the maintainers manually? While our research is still ongoing, we have found that providing guidance to developers with human-readable, actionable advice on how to fix quality and consistency issues is useful.



*Helm chart changerates in KubeApps Hub*

### Ecosystem components

Not only observe we commercial ecosystems, we also propose alternative improved designs. One success story was the launch of the Function Hub prototype in January 2018, which was the first online marketplace for cloud functions, more than one month before AWS launched their commercial SAR offering. The following screenshot gives an impression.



*Function Hub prototype by SPLab and PTI*

## Related software

- 🌀 **Functions Hub**. Open marketplace for deployable functions.
- 🌀 **FaaS-Timesharing**. Idleness optimisation in bag-of-tasks processing through cloud functions.
- 🌀 **Functions Collection**. Miscellaneous cloud functions deployable to FaaS.
- 🌀 **HelmQA**. Quality assurance for Helm charts and chart repositories.
- 🌀 **Snafu**. The Swiss Army Knife of Serverless Computing. (significant updates)

## Related talks

- 🌀 J. Spillner: “Serverless Delivery Hero — DevOps-style Tracing, Profiling and Auto-tuning of Cloud Functions“, Vienna Software Seminar, Vienna, Austria, December 2017
- 🌀 J. Spillner: “Serverless Computing: FaaSter, Better, Cheaper and More Pythonic“, 3rd Swiss Python Summit, Rapperswil, Switzerland, February 2018
- 🌀 J. Spillner: “Serverless Cyber-Physical Applications“ (keynote), Science Meets Industry, Dresden, Germany, March 2018
- 🌀 J. Spillner: “Helm Charts Quality Analysis“, Future Cloud Applications Meetup #5, Zurich, Switzerland, June 2018

## Related blog posts

- 🌀 “Extending Web Applications with Cloud Functions“, November 2017
- 🌀 “Portable Cloud Functions for Future-Proof Software Applications“, February 2018
- 🌀 “Cloud Function Marketplaces as Enablers of Serverless Computing Communities“, March 2018
- 🌀 “The role of FaaS in mixed-technology cloud and scientific computing applications“, June 2018

## Related tutorials

- 🌀 “Practical Tooling for Serverless Computing“, 10th International Conference on Utility and Cloud Computing (UCC), Austin, Texas, USA, December 2017. DOI: 10.1145/3147213.3149452
- 🌀 “Distributed Service Prototyping with Cloud Functions“, 38th IEEE International Conference on Distributed Computing Systems, Vienna, Austria, July 2018.

## / Service Tooling

However, Function Hub did not only convince by being earlier, but also by being better. Compared to AWS SAR, it accepts functions written for multiple cloud providers, not just AWS, as it is based on Snafu, our previous FaaS execution and prototyping tool. Furthermore, it is based on a federated XMPP infrastructure so that the entire catalogue content can easily be replicated, or functions can be tested at other nodes, leading to more flexibility and less vendor lock-in. We are ready to work further on the prototype and transfer it to any company aiming at being a major player in the serverless market.

## FaaS challenges

As the basic characteristics of commercial FaaS providers are now widely known and documented, along with appropriate function development methodologies, and open source tools are widely available, the research in the direction of FaaS becomes more specialised. One particularly interesting direction is the re-application of the time slice re-use which was also explored for virtual machines. Compared to the typical billing periods of one hour or one minute for machines, functions are microbilled in 100ms intervals, and therefore it is questionable if remaining idle periods per period can be re-used. It turns out that for functions actively pulling data and being able to control the amount of data to be pulled, this may be feasible; however, research is preliminary and no conclusive answer can be given to this problem.

In order to gain more insight into FaaS, we have produced several functions which are run at production-level in our lectures or tutorials where we cannot afford downtimes or degradation over time. Here, functions have already proven to be useful, in particular due to the occasional use pattern which allows for staying entirely within the free tier when the functions are hosted commercially.

## Researcher Spotlight: Piyush Harsh

Piyush Harsh is a senior 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 the H2020 Elastest project at a European level as well as in two Innosuisse-funded projects.



## Researcher Spotlight: Manu Perez Belmonte

Manu had been a Computer Engineering student from Mataró, Barcelona who first visited Switzerland while being in his final year of studies at Technocampus. He had joined the SPLab after the IAESTE internship program and started working as part-time assistant and master's student. He became very interested in learning new things, news technologies, meeting friends, sports and learning new languages. The assistant position had been the first study-related work experience for Manu as he was working in the CAB initiative team. He left SPLab in January 2018.



## Research Initiative: Cloud Accounting & Billing

### Project Spotlight: ElasTest



ElasTest is an EC-funded H2020 framework project that aims to simplify the management and provisioning of large scale tests for complex distributed systems. The project started on 1st Jan 2017 and will run until 31st Dec 2019.

SPLab's Cloud Accounting and Billing initiative is involved as task leaders for design and implementation of cost modeling and estimation engine. The objective of the task is quite significant as it will enable application developers not just estimate the cost of running their tests but in some way understand the monetary implications of executing their service as scale over diversified public / private cloud infrastructures.

In the first year of execution, we have designed the initial model to capture usage and cost information concerning various elements of test orchestration as well as application under observation. An early prototype has been released to the consortium and the effort is underway to integrate support services such as user-impersonation-as-a-service, browser-as-a-service, etc. with the first prototype.

CAB initiative aims to advance the initial model into a more comprehensive model with capability to capture fine grained feature set of a generalized test infrastructure, application under test as well as attach cost models that are close to what is being offered by public cloud providers today and in future. Such a model and the corresponding engine will enable fine grained analysis of applications under test, as well as deliver insights into putting such applications to operate at scale.

### Project Spotlight: Scale-UP

**swissuniversities** In SCALE-UP, which ended in the reporting period, a Distributed Computing Platform (DCP) in the cloud was developed that allows 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 needs, the DCP provides a matching algorithm that will guide to the most suitable one interactively. SCALE-UP was funded under a CUS programme by Swissuniversities.

## Researcher Spotlight: Oleksii Serhiienko

Oleksii Serhiienko is a part-time master student and researcher at the SPLab working on the CAB initiative. He has been graduated at the Kiev Polytechnic University majoring computer engineering and during his studies spent some time in Switzerland as research intern, working on OpenStack, in particular on the Ceilometer project. His code has then been added to the “Icehouse” OpenStack release.

Currently, Oleksii is working in the Amysta-SaaS project, identifying future architectures and middleware designs for multi-cloud management. He has contributed proof-of-concept code to the ManageIQ platform in this context.



## Researcher Spotlight: Serhii Dorodko

Serhii first joined the SPLab through the IAESTE internship program in 2016 while being a final year student in Computer engineering at National Technical University of Ukraine “Kyiv Polytechnic Institute”. His early work was within the Service Tooling initiative, conducting research in advanced cloud technologies such as FaaS. After graduation, he re-joined SPLab in the CAB initiative and now works on the COMBuST project to implement flexible pay-as-you-go microbilling on modern cloud platforms such as Kubernetes and OpenShift.



## / Cloud Accounting & Billing

### Project Spotlight: Amysta SaaS



Cloud computing is increasingly providing the resource and computing needs of businesses and enterprises worldwide. With well known companies adopting a cloud based operating roadmap, the confidence in the business propositions of cloud computing is on an upward trend. While many traditional sectors including banking, health, and public sector’s government services still do not wish to lose complete control on resource management, even they are increasingly adopting virtualization and private cloud deployments to improve utilization. And in between, there are numerous organizations that adopt a middle approach.

The common task undertaken by any CEO is optimizing the financial outlay by identifying how each division is using the budget, what is the cost of using the resources and whether there is a synergy among various departments in terms of resource usage or not. With the gradual shift from a fully in house infrastructure to outsourcing some of the computing needs to public providers of IaaS, PaaS, and SaaS, this task of cost management is increasingly becoming more challenging. The project aims to research models and algorithms addressing the concerns outlined above. Amysta SaaS is funded by Innosuisse and executed with Apalia.

### Project Spotlight: COMBuST



The modern software development process is undergoing a major paradigm shift with more organisations opting for DevOps software development and delivery. Modern cloud platforms have facilitated the transition to DevOps by supporting the creation of development environments, testing pipelines, staging and production rollouts on demand. This project will enhance microbilling frameworks by further enabling highly reactive accounting and billing runs in an ultra-dynamic environment of self-serviced, containerised workloads. The R&D work that will be conducted will enable sandboxed execution and analysis of various cost and billing models and understanding the impact of those on the company’s revenue baseline. The simulation toolkit will utilise anonymised user engagement patterns while executing a model in a sandboxed environment. This will enable the organisation to take a well informed decision before changing an established billing model to a new one. COMBuST is funded by Innosuisse and executed with VSHN AG.

## Related software

- ⚙️ **CMP-Testbed**. Recomputable evaluations of multi-cloud management platforms.
- ⚙️ **Sentinel**. Cloud application monitoring made easy.
- ⚙️ **Sentinel-Agents**. Sentinel monitoring agents.

## Related blog posts

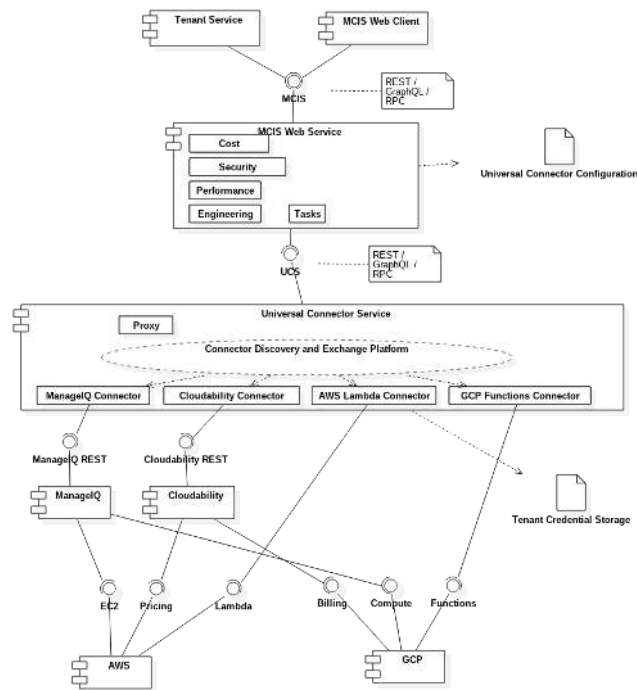
- ⚙️ “Cyclops 3.0 release with rule engine“, September 2017

## Related talks

- ⚙️ P. Harsh: “Sentinel: a monitoring platform for microservices“, Future Cloud Applications Meetup #4, Zurich, Switzerland, April 2018

## / Cloud Accounting & Billing

MULTI-CLOUD management has become an important topic for companies. Our work is concerned with designing a middleware around the concept of universal adapters which can not only connect to various cloud providers and stacks, but also get rich semantically annotated information out of each for the purpose of bottom-up information aggregation. The following figure shows an exemplary architecture which is still subject to further modification as our research continues.



*Universal connector architecture specification*

An interesting cross-link to our research on cloud functions has been identified when building connectors as stateless functions which are only enacted on demand. Compared to often monolithic and long-running management platforms, this will allow for greater flexibility in extending components at runtime and scaling out on demand. Furthermore, our work on portable (multi-)cloud functions contributes naturally to this extension. We intend to demonstrate this functionality in late 2018.


## Education and Qualification

AGAIN, we were involved in an almost habitual trip to another country to give guest lectures over the summer. And again, we had the usual autumn and spring lectures with bachelor and advanced studies students followed by the exams. And yet, our education profile sharpened compared to the previous year, and the first lab-associated student works contributed to that.

### Guest lectures in India



Upon invitation by the Department of Computer Science and Engineering of the Indian Institute of Technology (IIT) in Dhanbad, Jharkhand state in India, Josef Spillner went to give a talk, two seminars and a tutorial during two packed days in addition to attending a faculty council meeting and holding several bilateral meetings to discuss potential joint work. The visit yielded valuable exchange both on technical subjects and on desirable research and publication methods.



*Lecturing inside IIT Dhanbad*

⚙️ Blog post: “Guest Lectures at Indian Institute of Technology – IIT (ISM) Dhanbad“, August 2018

### Student works

Computer science student Fabian Camenzind conducted his bachelor thesis titled «Migrating with microservices to cloud-native application architectures» in conjunction with local IT company ti&m AG. The thesis was defended successfully and opens up a much wider perspective on such applications, considering also frontends and development aspects in addition to the runtime. Part-time master student Oleksii Serhienko, who also works in the lab as assistant, successfully concluded his practical modules «Survey on Re\* research» and «Recomputable testbed for API abstraction layers and connectors», the latter being submitted as paper to CloudCom'18.

## /Education and Qualification

### Internet Service Prototyping (bachelor, elective module)

Our favourite elective module, Internet Service Prototyping, took place in autumn 2017 and culminated in team presentations in December. Several teams presented distributed cloud-enabled software applications, sometimes even multi-cloud capable, which can be microbilled for on-demand use. Some brave students even explored dual-technology applications, combining cloud functions with containers.

### Scripting (advanced studies module)

Now happening every semester due to the high interest in data science, our Scripting module is filled with 30 business and public administration employees in every iteration. The students explore supervised and unsupervised machine learning tasks after spending most of the time learning the basics of scripting with Python. We support the e-didactical aspect of teaching with serverless web applications, causing few maintenance issues while providing powerful features such as score and rank calculation as part of a gamified learning environment.

⚙️ Blog post: “Probabilistically Stateful Serverless Web Applications“, October 2017

### Doctoral theses

SPLab still co-supervises two doctorate propositions in its areas of expertise: Stanley Ramalho Lima at University of Coimbra, Portugal, and Ambrósio Patrício Vumo, at Universidade Eduardo Mondlane, Mozambique. Ambrósio spent some time in our laboratory in October 2017 along with Martina Barros in order to improve his understanding of OpenStack, learn about local cloud providers and privacy regulations, and dive deeper into research. He has since continued his field research in Mozambique and worked on the next publication to define a cloud security framework suitable for this country.



*Visitors from Mozambique*

⚙️ Blog post: “Guest Researchers from Mozambique“, October 2017

## Lab Life and Qualification

LOOKING beyond the daily research work is of high importance to the long-term fitness of our laboratory. We therefore afford to go the extra mile and foster better research environments.

### Research and publishing approach

After exploring in depth preprints and open access strategies in the year before, we now concentrated on reproducibility and open proceedings. In particular, Oleksii Serhiienko wrote a study piece titled «Survey on Re\* Research» which compares repeatability, reproducibility, recomputability, resusability and replicability, and Josef Spillner experimented with open reviews. Eventually, the goal is to produce a lab-wide guideline on how to properly conduct and communicate research.

### International exchange

Following up on the teaching-centric trip of Josef Spillner to Itaipu Technology Park in Paraguay in August 2017, from January to March 2018 two researchers from PTI, Yessica Bogado and Walter Benítez, visited the Service Prototyping Lab to discuss research directions and ideas.

⚙️ Blog post: “Visitors from Itaipu Technology Park“, January 2018



*Visitors from Paraguay*

We further intensified international relations by visiting AGH University of Science and Technology in Kraków, Poland, as well as Indian Institute of Technology (IIT) in Dhanbad, India.

### SPLab alumni

Serhii Dorodko concluded his internship at SPLab in August 2017, proceeded to finish his master's theses at the National Technical University of Ukraine (Igor Sikorski Kyiv Polytechnic Institute), and re-joined SPLab as scientific assistant in April 2018. Martin Skoviera left SPLab in September 2017, Manuel Perez Belmonte in January 2018, Manuel Ramírez López in June 2018 and Stefan Junker in July 2018, all four remaining as well-qualified software developers in the greater Zurich area to strengthen the local ICT industry.



## Events

### Presence at events

In recent months covered by the reporting period, researchers from the Service Prototyping Lab participated in the following events with keynotes, talks, tutorials and technical demonstrations.

- 6th European Conference on Service-Oriented and Cloud Computing (ESOCC), Oslo, Norway, September 2017.
- 10th IEEE/ACM International Conference on Utility and Cloud Computing (UCC), Austin, Texas, USA, December 2017.
- Vienna Software Seminar, Vienna, Austria, December 2017.
- 3rd Swiss Python Summit, Rapperswil, Switzerland, February 2018.
- Keynote @ Science Meets Industry, Dresden, Germany, March 2018.
- 8th International Conference on Cloud Computing and Services Science (CLOSER), Funchal - Madeira, Portugal, March 2018.
- 4th IEEE International Conference on Recent Advances in Information Technology (RAIT), Dhanbad, India, March 2018.
- Open Cloud Day, Winterthur, Switzerland, May 2018.
- EclipseCon, Toulouse, France, June 2018.
- 38th IEEE International Conference on Distributed Computing Systems (ICDCS), Vienna, Austria, July 2018.



*Impressions from UCC 2017 held in the Lone Star state Texas*

- ⚙️ Blog post: “ESOCC 2017 – Oslo“, October 2017
- ⚙️ Blog post: “UCC 2017 Coverage – Day 1“, December 2017
- ⚙️ Blog post: “Vienna Software Seminar: DevOps and Continuous-\*\*\*“, December 2017
- ⚙️ Blog post: “Impressions from Swiss Python Summit 2018“, February 2018
- ⚙️ Blog post: “Science Meets Industry and Innovation Alignment“, March 2018
- ⚙️ Blog post: “CLOSER’18 conference report“, March 2018
- ⚙️ Blog post: “Open Cloud Day 2018“, June 2018
- ⚙️ Blog post: “Brief report on the ICDCS’18 conference“, July 2018
- ⚙️ Blog post: “EclipseCon’18 France – Recap“, July 2018

## Community

### Academic services

Researchers from the Service Prototyping Lab have filled several volunteer roles in the reporting period to drive journals and magazines.

- Guest Editor, IEEE Cloud Computing Magazine (Josef Spillner)
- Guest Editor, Springer Empirical Software Engineering Journal - EMSE (Sebastiano Panichella)
- Guest Editor, Elsevier Information and Software Technology - IST (Sebastiano Panichella)
- Editorial Board Member, Wiley Journal of Software: Evolution and Process (Sebastiano Panichella)

Further academic service duties have been accomplished in relation to conferences and workshops.

- General Chair, 11th IEEE/ACM UCC and 5th IEEE/ACM BDCAT 2018 (Josef Spillner)
- Demo and Poster Chair, 19th ACM/IFIP/USENIX MIDDLEWARE 2018 (Piyush Harsh)
- Workshop Chair, Cloud Software Engineering at IEEE CLOUD 2018 (Josef Spillner)
- Workshop Chair, International Workshop on Serverless Computing – WoSC (Stefan Junker)
- Workshop Chair, Cloud-Native Applications Design and Experience - CNAX (Sebastiano Panichella)
- Scientific Track Chair, European Symposium on Serverless Computing and Applications – ESSCA (Josef Spillner)

As most of these duties are far from over, we are now looking forward to a successful materialisation of the events. Additionally, reviewer duties and technical programme committee memberships were concluded for the following venues.

- Journals: ACM TOIT, ACM CSUR, IEEE TCC, Springer JIIS
- Magazines: IEEE CCM
- Conferences: ACM MEDES'17, IEEE CloudCom'17, IEEE/ACM UCC'17, IEEE/ACM BDCAT'17, INSTICC CLOSER'18, IEEE NetSoft'18, EUSPN'18
- Workshops: CloudAM'17, QUDOS'18, DSS'18

Open reviews were moreover tried to complement attempts for modern research and publication approaches.

### Local events

We have again organised two meetups/evening seminars, Future Cloud Applications #4 and #5, on various topics around our research initiatives. Subscribe to the meetup group online: <https://www.meetup.com/Future-Cloud-Applications/>

Furthermore, we continue to operate and moderate the cloud-announce mailing list on behalf of the local and global cloud computing research community.

## Summary of Research Outputs

As in previous years, we adopt a broad definition of research outputs in line with recent trends at funding bodies. The outputs encompass reviewed and non-reviewed publications, datasets, testbeds and research software.

### Peer-reviewed publications (4)

- ⚙️ M. Ramírez López, J. Spillner: “Towards Quantifiable Boundaries for Elastic Horizontal Scaling of Microservices“, 6th International Workshop on Clouds and (eScience) Applications Management (CloudAM) / 10th International Conference on Utility and Cloud Computing Companion (UCC), Austin, Texas, USA, December 2017. DOI: 10.1145/3147234.3148111
- ⚙️ G. Toffetti, T. Lötscher, S. Kenzhegulov, J. Spillner, T. M. Bohnert: “Cloud Robotics: SLAM and Autonomous Exploration on PaaS“, 6th International Workshop on Clouds and (eScience) Applications Management (CloudAM) / 10th International Conference on Utility and Cloud Computing Companion (UCC), Austin, Texas, USA, December 2017. DOI: 10.1145/3147234.3148100
- ⚙️ J. Spillner, Y. Bogado, W. Benítez, F. López Pires: “Co-transformation to cloud-native applications: Development experiences and experimental evaluation“, 8th International Conference on Cloud Computing and Services Science (CLOSER), Funchal - Madeira, Portugal, March 2018. DOI: 10.21256/zhaw-3508
- ⚙️ C. V. Alexandru, J. J. Merchante, S. Panichella, S. Proksch, H. Gall, G. Robles: “On the Usage of ‘Pythonic’ Idioms in Python Programs“, Onward!, Boston, Massachusetts, United States, November 2018. To appear.

### Preprints (1)

- ⚙️ P. Leitner, E. Wittern, J. Spillner, W. Hummer: “A mixed-method empirical study of Function-as-a-Service software development in industrial practice“, PeerJ Preprints 6:e27005v1, June 2018.

### Data repositories and testbeds (7)

- ⚙️ O. Serhienko, J. Spillner: “Cloud management platform evaluation data generated by CMP<sup>2</sup>“
- ⚙️ P. Leitner, E. Wittern, J. Spillner, W. Hummer: “Survey and Interview Data from Mixed-Method Survey of Serverless Computing and Function-as-a-Service Software Development in Industrial Practice“
- ⚙️ J. Spillner: “FaaS Characteristics and Constraints Knowledge Base“
- ⚙️ J. Spillner: “Serverless Literature Dataset“
- ⚙️ J. Spillner: “Duplicate reduction in Helm charts“
- ⚙️ J. Spillner: “Docker Resilience“
- ⚙️ J. Spillner, M. Ramírez López: “Quantifiable Scaling“