

SRT-15

Intelligence Push in the Enterprise Realm

Project Number:	FP7-ICT-257843
Project Title:	SRT-15: Intelligence Push in the Enterprise Realm
Deliverable Number:	D5.4
Title of Deliverable:	Research Newsletter 1
Contractual Date of Delivery:	28/09/2011
Actual Date of Delivery:	28/09/2011

Abstract

This deliverable is a public research newsletter that presents the objectives, activities, achievements and events of the project to a general audience.

List of Contributors

Name	Organization	E-mail
Etienne Rivière	UniNE	Etienne.riviere@unine.ch
Pascal Felber	UniNE	Pascal.felber@unine.ch
Jean-François Pineau	UniNE	Jean-francois.pineau@unine.ch
Emanuel Onica	UniNE	Emanuel.onica@unine.ch
Raphaël Barazzutti	UniNE	Raphael.barazzutti@unine.ch
Luigi Romano	Epsilon	Luigi.romano@epsilononline.com
Luigi Sgaglione	Epsilon	Luigi.sgaglione@epsilononline.com
Luigi Coppolino	Epsilon	Luigi.coppolino@epsilononline.com
Salvatore D'Antonio	Epsilon	Salvatore.dantonio@epsilononline.com
Rosario Cristaldi	Epsilon	Rosario.cristaldi@epsilononline.com
Antonio Sergio	Epsilon	Antonio.sergio@epsilononline.com
Christof Fetzter	TU Dresden	christof.fetzter@se.inf.tu-dresden.de
André Martin	TU Dresden	andre.martin@se.inf.tu-dresden.de
Stefan Weigert	TU Dresden	stefan.weigert@se.inf.tu-dresden.de
Marco Serafini	Yahoo! Iberia	serafini@yahoo-inc.com
Flavio Junqueira	Yahoo! Iberia	fpj@yahoo-inc.com
Matthieu Morel	Yahoo! Iberia	matthieu@yahoo-inc.com
Daniel Gómez Ferro	Yahoo! Iberia	danielgf@yahoo-inc.com
Berkant Barla Cambazoglu	Yahoo! Iberia	barla@yahoo-inc.com
Zbigniew Jerzak	SAP	Zbigniew.Jerzak@sap.com
Thomas Heinze	SAP	Thomas.Heinze@sap.com
Sandra Steinbrecher	SAP	Sandra.Steinbrecher@sap.com

Document Approval

	Name	Email	Date
Approved by WP Leader	Luigi Romano	luigi.romano@epsilononline.com	2011.10.20
Approved by GA Member 1	Zbigniew Jerzak	zbigniew.jerzak@sap.com	2011.10.21
Approved by GA Member 2	Christof Fetzer	christof.fetzer@se.inf.tu-dresden.de	2011.10.21

Contents

LIST OF CONTRIBUTORS II

DOCUMENT APPROVAL.....III

CONTENTS..... IV

LIST OF FIGURES V

List of Figures

Figure 1: SRT-15 platform2

Figure 2: SRT-15 project members in Champéry, Switzerland. March 2011.3

Figure 3: Partners of the project and their geographical situation3

Figure 4: architecture of the small-scale modular CBR engine4

Figure 5: Overall architecture of SRT-15 Use Cases.....7

Figure 6: design and runtime of the SRT-15 DEBS Challenge Prototype..... 10

Figure 7: SRT-15 project member André Martin (TUD) presenting the prototype at the J.T. Watson
Research Center in Yorktown, NY..... 10

GLOSSARY

SRT-15 research newsletter, issue 1.
Editor: Dr. Etienne Riviere, University of Neuchâtel

In this issue

1. Editorial: a word from the project coordinator
2. The SRT-15 consortium
3. Update on SRT-15 technologies and achievements
 - a. Content-Based Routing engine
 - b. Complex Event Processing engine
 - c. Dependability
 - d. Privacy
 - e. SRT-15 applications and uses cases
4. Highlight: SRT-15 makes it to the final of the DEBS 2011 Grand Challenge!
5. Collaboration and dissemination efforts
6. Dissemination highlight: SRT-15 presented with a video thanks to the HOLA! project



Editorial - a word from the project coordinator: The shape of enterprise or institutional networks and IT systems in general is rapidly changing, with a notable increase in scale, complexity and inter-dependency. An answer to the ever-increasing complexity, diversity and distributed nature of computing infrastructures is the use of cloud computing: outsourced computing resources, provided on a pay-as-you-go fashion by external providers. Cloud services allow an easy scaling and deployment of multi-locations and even multi-companies IT systems, without the need for user companies to provision in advance necessary but costly resources. Cloud computing implies outsourcing and management of resources by third parties. The elasticity of resource allocation and their shared nature as well as the complexity and scale of information flows in large companies calls for appropriate platforms and enabling technologies that can foster the uptake of the promises of cloud computing. Two other aspects are of paramount importance to support this uptake. First, in large-scale cloud environments faults are the norm rather than the exception: server components fail or corrupt information and messages are being lost due to misconnections. The uptake of cloud technologies will require the provision to the users of strong guarantees on dependability of the proposed platforms and technologies. Second, the shared nature of public cloud offers typically lead to suspiciousness for companies' IT specialists: they distrust cloud providers for completely protecting the privacy and safety of their data. Moreover there exist almost no means to control these aspects – with recent reports on co-location attacks in public clouds only proving them right. As a result, the sound solution is to enable privacy-preservation as a premier feature of a cloud system and not as an additional one. Guarantees to the client on the privacy of the data and information flow are a key for a better spread of cloud usage in companies.

The aforementioned goals drive the SRT-15 project. This project started on October 2010 and is being funded by the EC under the FP7 programme for a period of 30 months. In order to be able to embrace the change in the enterprise information processing landscape, SRT-15 relies on technologies that support rapid change: cloud computing, content-based routing, and complex event processing. SRT-15 not only embraces change, it also controls it. Privacy and dependability have been recently identified as main challenges associated with the on-demand model. Therefore, SRT-15 will by design support extreme robustness and will preserve the privacy of data, which is processed in the public and private clouds SRT-15 is deployed upon. This will allow SRT-15 to maintain control over the dependability and privacy of the enterprise data despite failures and breaches of the underlying infrastructure.

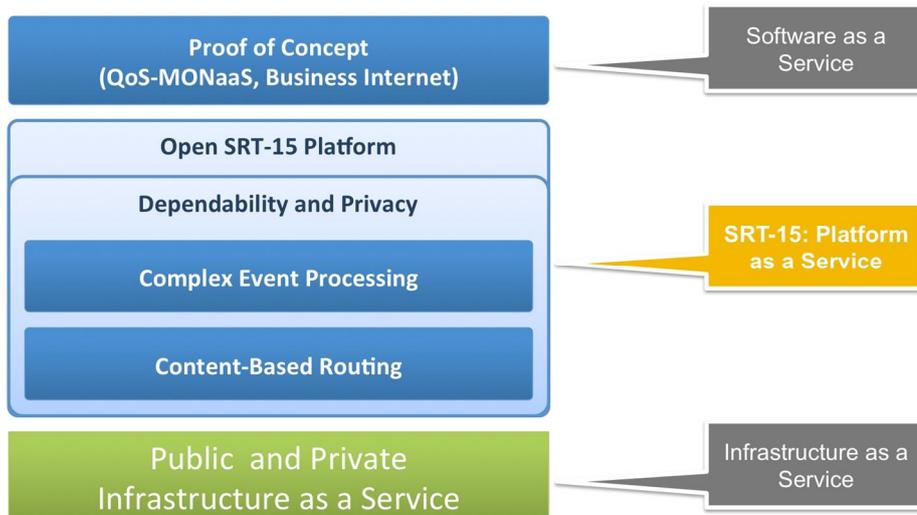


Figure 1: SRT-15 platform

The SRT-15 project is organized around three main objectives, which are the definition of the complex event processing and content-based routing technologies, and the exploration and definition of dependability, privacy and scalability mechanisms. Figure 1 presents a schematic view of the SRT-15 components, with proof of concept applications on top and supported infrastructures as a service below. This first research newsletter of the SRT-15 project presents some early achievements made by members of the project towards these goals. We present short summaries of the work done in each work package, highlight the positioning of the project's solutions against state-of-the-art systems, and our successful participation to the DEBS 2011 challenge. Finally, we list upcoming events and milestones for the project and current and future collaboration efforts.

The SRT-15 consortium



Figure 2: SRT-15 project members in Champéry, Switzerland. March 2011.

The SRT-15 project is composed of three industrial partners: SAP AG (coordinator of the project), Yahoo! Iberia and Epsilon, and two academic partners: Technische Universität Dresden and University of Neuchâtel.



Figure 3: Partners of the project and their geographical situation

SRT-15 technologies: Content-based routing

The effort of SRT-15 project members on content-based routing (CBR) technologies has been mostly concentrated in WP1 during this initial part of the project. We investigated mechanisms to support efficient and modular filtering mechanisms, with the requirements for scalability and dependability in mind from the beginning of the project. The CBR mechanisms of SRT-15 rely on the interplay with the complex event processing (CEP) engine, and benefits from the guarantees that are, or will be offered by this engine (scalability, dependability, availability to name a few). We concentrated our initial efforts towards the construction of a single-node matching engine that allows extreme throughput and flexibility. The idea is to later evolve this simple but well-principled design to support the scalability and security needs of the project, which will be the main topics of interest for the next period. In particular, this prototype already allows content-based publish and subscribe to happen between clients using different languages and hardware architectures, and is designed to support heterogeneity both for clients and for other matching engines that will be instantiated in the sequel of the project to support scalability.

One of the main components of a CBR engine is its filtering algorithm, which we implement by means of modular filtering libraries. Its goal is to match incoming publications against stored subscriptions, and to determine which subscribers shall receive each publication. We have made progress on the matter by designing efficient filtering techniques that can aggregate subscriptions in efficient filtering trees, leading to sub-linear filtering complexity. We also conducted a study on the scalability model and mechanisms, where matching engines are being replicated and subscriptions and publications partitioned to achieve a natural scalability capability, which drove a large part of the design of the single-node CBR engine. Part of this work concerned the use of clustering techniques for partitioning subscriptions in an efficient manner w.r.t. the filtering operation itself. Finally, initial work on supporting privacy-preservation has been conducted, by exploring the various algorithms available and proposing detailed analysis and improvements to the one selected. We observed in particular, that the support of privacy-preservation induces that coverage relations between subscriptions can no longer be used to support efficient filtering algorithms; hence, we started developing strategies for efficiently pre-filtering publications before running the actual filtering based on encrypted content. All these mechanisms are envisioned as the components of the next step of the CBR engine, which will support scalability in term of clients and number of subscriptions.

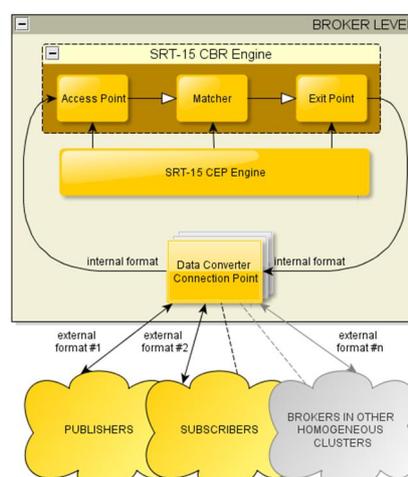


Figure 4: architecture of the small-scale modular CBR engine

SRT-15 technologies: Complex Event Processing

During the initial phase of the project, SRT-15 project members working on Complex Event Processing have been focusing on analysing the requirements for the CEP engine with regards to the use cases of SRT-15 and designing the base architecture of the CEP engine. The analysis covered identification of required operators and investigation of aspects such as the data and event model, the temporal model, the query model and the system model based on related work.

To support heterogeneity between the machines and clouds onto which the CEP engine will be deployed, we use GoogleProtocol buffers as our event format to exchange messages in heterogeneous environments across different architectures. Furthermore, the CEP engine will offer six parameterizable standard CEP operators such as selection, join, negation, projection etc. to SRT-15 users in addition to an interface which allows users to define their own custom operators.

In the second phase of this first year of the project, a first prototype of the CEP engine has been implemented on top of the mature code base of StreamMine [SM]. For the implementation of operators, a state machine based approach has been chosen. Therefore, standard as well as custom written operators are expressed as state machines using a domain specific language (DSL). The DSL allows the translation of custom written operators into highly optimized C/C++ code by the open source SMC compiler [SMC].

In addition to the implementation of the CEP engine, the applicability of the approach has been assessed by participating in the DEBS challenge (see the dedicated highlight in this newsletter for more details).

Finally, a set of micro-benchmarks has been performed to identify performance issues as early as possible, which will be reflected in design and implementation decisions for the following phases. The benchmarks included throughput and latency measurements with regards to scalability horizontal as well vertical wise as well as the impact of the operator layout due to the graph based design.

In the next phases, an extension of the prototype is anticipated to support elasticity as well as dynamic operator placement.

SRT-15 technologies: Dependability

The SRT-15 services run on top of the Complex Event Processing system. The dependability work on SRT-15 focused on enhancing the fault tolerance of CEP, thus enhancing the dependability of all layers on top of it including CBR. SRT-15 improves on the state of the art in dependability in two topics.

First, it proposes a new approach to the tolerance of data corruptions that endanger the integrity of data and computation. Today, the only ways to protect distributed systems against data corruptions in a systematic manner, that is, with formal dependability guarantees, are Byzantine-fault tolerance and coding. Byzantine-fault tolerance is a method to tolerate worst-case faults, including intrusions. In cloud platforms, Byzantine fault tolerance is considered overkill because it also covers security issues that, in modern cloud systems, are handled by orthogonal techniques owing to the field of security rather than dependability. Furthermore, while there are known ways to make simple client-server system Byzantine fault tolerant using existing libraries, designing new complex distributed systems that are Byzantine-fault tolerant is very complex endeavour that is often subject of PhD theses. Alternatively, one can use coding techniques, but these can only limitedly protect data while is being

computed (for example, they do not tolerate the corruption of the program counter while an instruction block is being executed) and often require compiler support to be automated (like SWIFT). In SRT-15 we propose a new fault model, called Arbitrary State Corruptions (ASC), which is more suited to the needs of dependable cloud services because (a) by restricting to state corruptions, it does not have to pay the complexity costs of tolerating intrusions and handling other security matters that are orthogonal to dependability, (b) it allows defining a generic transformation technique that can be applied to any crash tolerant distributed system to harden it against ASC faults, and (c) it covers a wider range of faults than any existing coding technique, and it provides end-to-end provable guarantees. A first application of the transformation technique was hardening a state machine replication algorithm based on Paxos. Initial evaluation shows that the approach has limited cost.

The second improvement of SRT-15 compared to the state of the art in dependability is a technique for efficient checkpointing techniques for CEP and stream processing systems in general. SRT-15 supports precise recovery from failures, ensuring that faults have no effect whatsoever on computation. Ensuring this very strong fault tolerance guarantee is typically quite complex and expensive in terms of performance overhead. SRT-15 proposes a new approach based on virtual synchrony that allows every CEP process to take checkpoints in isolation, without prior coordination with other processes. Experimental evaluation shows that this form of checkpointing has minimal performance overhead.

SRT-15 technologies: Privacy

The SRT-15 platform will support privacy preservation at the level of the CBR engine. Indeed, some clouds may be given a lower trust than others, especially when they are provided by a third party and potentially shared with other customers. We call them public clouds in this case. An example is Amazon's EC2, which is a well-known IaaS cloud offer. Recent evidence shows that the shared nature of such public clouds could put the data manipulated by the machines on a public cloud vulnerable to eavesdropping [RTSS09]. The use of virtualization could allow an attacker to co-locate its virtual machines on the same physical machine as the one she wishes to observe. This observation is then made possible by exploiting potential vulnerabilities and bugs in the hypervisor and operating systems. Another scenario is the presence of a malevolent operator at the cloud provider, that could wish to obtain critical information from the clients of the cloud operator, e.g., for business intelligence gathering about ones' competitors.

There is however a mismatch between the protection of privacy of content in public clouds and the use of content-based routing in these clouds. How can one decide on the destination(s) of a publication based on its content, if this content is not accessible (e.g., encrypted by a key that is not available to the public cloud machines)? Moreover, the subscriptions themselves can reveal important information about the nature of the application and the business logic behind the data flow of an application: for instance, in a stock trade application, subscriptions can reveal interest for some stock and thus the investment strategy of a given financial company. For these two reasons, there is a clear necessity to protect the privacy of the content exchanged, and protecting the privacy of the interest for content (respectively, publication- and subscription-privacy).

The approach followed by the SRT-15 project, and investigated both in WP3 (for the theoretical part) and in WP1 (for the implementation part) is to use privacy-preserving routing mechanisms, which is a special form of encrypted processing. The idea is to encrypt both subscriptions and publications in a way that still allows to perform the matching operation: determining if one publication matches the interest of one subscription, but preventing from performing operations such as range coverage between subscriptions, aggregation of subscriptions, or clustering of subscriptions based on their ex-

pressed interests, all of which that would allow inferring the subscription based on a greedy discovery algorithm.

The algorithm of choice of SRT-15 privacy-preserving CBR will be ASPE [CGB10], which is itself based on scalar-product-preserving transformation. This algorithm, its performance (theoretical and practical) will be further analyzed and some improvements on speed and efficiency are already being developed in the context of the project. These will be integrated in the final platform to support seamless integration of regular and privacy-preserving CBR depending on the level of trust granted by one application to the various clouds on which it is executing. In order for applications to declare the trust they are granting to the different clouds, we proposed a trust model that is based on the notion of trust properties, allowing fine-grain control over what operation can, or cannot, be performed on plaintext data on each type of clouds. This trust model is used to support the automatic key exchange between clouds to support decryption or encryption at the clouds borders.

SRT-15 applications and use cases

Two proof of concept applications are being developed to create a convincing, integrated proof of concept which will increase the dissemination and impact potentials of the SRT-15 project, namely:

- Smart Meters: a variant of an Internet of Things (IoT) application developed by SAP AG, which implements remote monitoring of power consumption in a Smart Grid environment
- QoS-MONaaS (Quality of Service MONitoring as a Service): a dependable QoS monitoring facility developed by EPSILON which relies on the "as a Service" paradigm, and is thus made available to virtually all SRT-15 platform users in a seamless way.

The high-level scenario overview is presented in Figure 5.

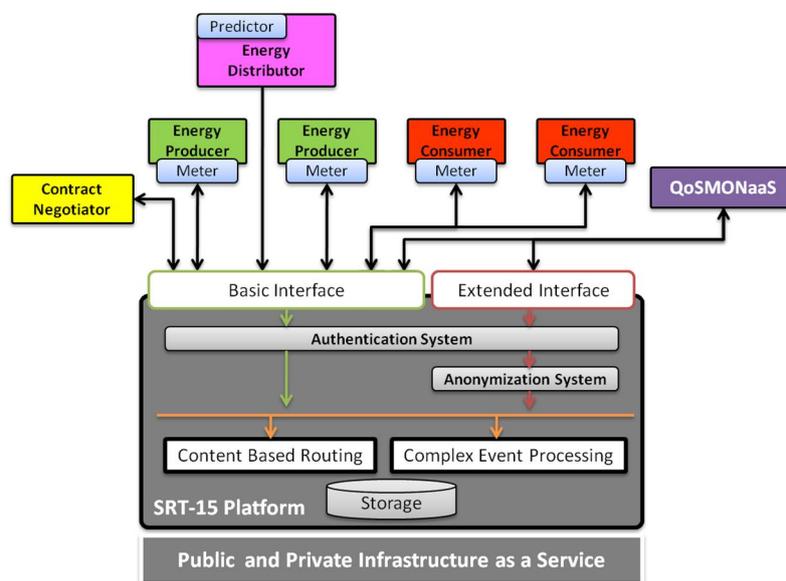


Figure 5: Overall architecture of SRT-15 Use Cases

Smart Meters uses the SRT-15 platform for distribution and real-time processing of measurements related to energy consumption and production. The Energy Distributor sells power to Energy Consumers, based on a set of pre-defined SLAs. To do so, it buys energy from Energy Producers, also based on a set of pre-defined SLAs. Energy Producers and Energy Consumers are equipped with digi-

tal devices that make measurements on actual production/consumption available to the SRT-15 platform. In the existing setup, energy measurements (for both production and consumption) are collected every 15 minutes. Even though the physical devices are capable of producing measurements much more often, it is currently being discussed to provide only aggregated information, mostly due to privacy reasons [TKG11]. In a traditional (i.e. not based on cloud computing technology) set-up, energy measurements are encapsulated in form of events and are propagated “upwards” towards a measurement collection point intermediary. A hierarchy of intermediaries gathers energy data (events) from multiple meters thus collecting energy consumption information for a certain area. Such a structure can be hierarchically built up to cover larger and larger areas. In a cloud-based environment, Energy Producers and Energy Consumers data is forwarded for processing in the cloud. The Energy Distributor - continuously, using the SRT-15 Basic Interface - gathers the data on power consumption, and processes it in an attempt to predict the energy demand for the next time frame. The objective is to allow the Energy Distributor to efficiently manage renewable energy and balance supply and demand. Measurements are also persisted in the SRT-15 storage component. The storage component serves as input to additional components, which do off-line data analyses on the historical. Based on the result of the prediction process, the Energy Distributor knows that it needs to either contact additional Energy Producers and add them to its pool or it may get rid of some of the Energy Producers which are in the current pool. Contracts between parties are handled by the Contract Negotiator. The abovementioned actors (namely: Energy Distributor, Energy Consumers, Energy Producer, and Contract Negotiator) collectively form the Smart Meters application. An experimental testbed has been set up, where all parties of the Smart Meters application utilize the cloud-based routing and processing capabilities of a preliminary version of SRT-15 platform, and QoS-MONaaS is in charge of monitoring the actual QoS which is actually delivered to individual parties.

More details on the current implementation of SRT-15 use case architecture and preliminary experimental results can be found in [RDJF11] and [CCCDR11].

Highlight: SRT-15 makes it to the final of the DEBS 2011 Grand Challenge!

Members of the SRT-15 participated to the 2011 DEBS Grand Challenge [DGC], organized in the context of the DEBS 2011 conference (Fifth ACM International Conference on Distributed Event-Based Systems) in New York City, USA. The contributions to the grand challenges were presented during the conference on July 11-15 2011.

The objective of the DEBS Grand Challenge is to promote the research and development in the area of event processing by enabling R&D groups (research and industry ones) and software vendors to demonstrate the strength of their results.

In this year's competition, participating teams were challenged to implement an ESP system for a *Trivia Geeks Club* game as follows: The system receives different types of events continuously such as question, answer and answer annulment events from participating players. Events must then be evaluated using a predefined set of rules to keep the player's score up to date. The majority of such rules require the recognition of simple event patterns such as the following: "three answers incorrect without a correct answer in the middle = -50pts". A common approach for pattern recognition is the utilization of state machines.

The SRT-15 research prototype consists of the following components as shown in Figure 6. A Run-Time equipped with a high performance CEP engine and a Design-Time which allows SRT-15 users to model CEP operators (expressible as state machines) quickly at different levels of abstracts: User can either design their operators in a visual way via drag and drop through the tight integration of the open source tool RapidMiner [RM], or alternatively at a lower level using the DSL of the open source SMC (state machine compiler) [SMC]. Once operators are defined, the run time of the SRT-15 research prototype transforms those operators in highly optimized C/C++ code and furthermore deploys them onto the cloud. Furthermore, the Design-Time offers offline debugging of applications prior deploying them onto the cloud.

SRT-15 research prototype has been chosen as one of the finalist for the DEBS Grand Challenge 2011 among ETH Zürich and the Esper-Tech Team. The solution was presented during the DEBS 2011 at J.T. Watson Research Center, Yorktown, NY and was recognized as an innovative and intuitive to use solution.

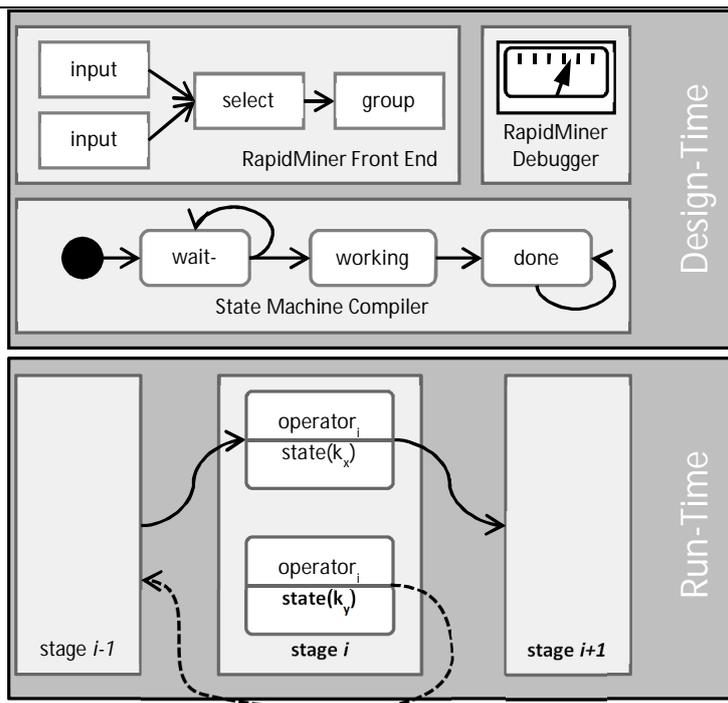


Figure 6: design and runtime of the SRT-15 DEBS Challenge Prototype



Figure 7: SRT-15 project member André Martin (TUD) presenting the prototype at the J.T. Watson Research Center in Yorktown, NY

Collaboration and dissemination efforts

Members of the SRT-15 project are active in establishing successful collaboration plans with various other projects, funded by the European Commission and other funding agencies.

After investigating a number of potential collaboration partners in the deliverable "D5.2 - Collaboration Plan" it has been decided to pursue technical collaboration with a small number of very well aligned projects. In particular, the SRT-15 research project has identified opportunities for cooperation with the following projects funded within the Objective 1.2 "Internet of Services, Software and Virtualisation" and "Service and Software Architectures, Infrastructures and Engineering" of the FP7-ICT Programme for Research and Development:

- PLAY (<http://www.play-project.eu/>), which aims to "*develop and validate an elastic and reliable architecture for dynamic and complex, event-driven interaction in large highly distributed and heterogeneous service systems*". The cooperation activities include regular exchange with PLAY partners (Nenad Stojanovic and Roland Stühmer from the FZI Forschungszentrum Informatik an der Universität Karlsruhe) responsible for the development of the inference-based Distributed Complex Event Processing engine. Future collaboration plans encompass joint participation in a planned EuroSys 2012 workshop organized by the SRT-15 as well as possible joint participation in the summer school organized by PLAY in 2013.
- MoSAIC (<http://www.mosaic-cloud.eu/>), which aims to "build an application platform that negotiates cloud services as requested by their users".

We are currently identifying synergies and collaboration plans with the members of these projects. In particular, in conjunction with the SRT-15 plenary meeting in Procida (Naples), a technical meeting with representatives of the MoSAIC project is planned, to foster the cooperation between the two projects.

SRT-15 representatives have attended several dissemination/collaboration events organized by the European Commission. In particular the 2010 edition of the Internet of Services [IoS2010] and the 2010 and 2011 editions of the Future Internet Assembly [FIAGhent] [FIABud].

Next planned events are:

- Internet of Services 2011 - Collaboration meeting 28-29 September 2011, Diamant Center Brussels, where SRT-15 will be presented in following sessions:
 1. Trust and Security
 2. Real-Time Data Management and Processing
 3. Virtualized Service Platforms
 4. Semantics in Services and Cloud
 5. QoS & SLAs
 6. Service Engineering
- FIA Poznan 25-26 October (to be confirmed)

Also, issues related to QoS Monitoring in cloud environments were discussed with representatives of projects participating to the "Cloud Computing: Projects and Initiatives" workshop which was organized within the context of the EUROPAR 2011 International Conference, held on August 30 2011 in Bordeaux. On that occasion, the SRT-15 approach was also presented to the audience.

On the 22nd of August 2011 SAP organized a SRT-15/STRATOSPHERE workshop in Dresden. The goal of the workshop was to bring together research groups working on SRT-15 research project and the STRATOSPHERE project (<http://www.stratosphere.eu/>). Key speakers included Prof. Christof Fetzer from the TU Dresden as well as Prof. Volker Markl from the TU Berlin. STRATOSPHERE is a joint project between the Technische Universität Berlin, Humboldt Universität zu Berlin, and the Hasso-Plattner-Institut in Potsdam and is funded by the German Research Foundation (DFG - <http://www.dfg.de/>). The goal of STRATOSPHERE is to research topics revolving around the "Information Management in the Cloud" and specifically to "explore the power of massively parallel com-

puting for complex information management applications". Prof. Volker Markl acts as speaker of the STRATOSPHERE unit. The workshop was attended by fifteen researchers and served as a platform for discussion and dissemination of SRT-15 and STRATOSPHERE related ideas. It included four keynote presentations and two system demonstrations, including the presentation of the SRT-15 prototype presented at the DEBS Challenge. The goal of the SRT-15 demonstration was to provide feedback for the platform evaluation from external users.

Dissemination highlight: SRT-15 presented with a video thanks to the HOLA! project

The HOLA! Project [HOLA] ran a contest (HOLA! Video Contest) to select the best 3 European R&D projects related to the Internet of Services with great potential impact on society and EU economy. Selection criteria were:

1. Potential impact of project's results on its end users
2. Potential importance of project's results for EU society
3. Potential importance of project's results for EU economy
4. Potential strength of the effect of the clip for dissemination of project's results
5. General evaluation of the application.

SRT-15 has been selected as one of the three winners of the HOLA! Video Contest (the others two winners being STREAM [STREAM] and PLAY [PLAY]). The SRT-15 staff is currently working with the HOLA! team to the production of a promotional video, which will be developed and financed by HOLA!. The video highlights the project motivation and features, for dissemination and exploitation purposes. The video clip will be presented at the IoS Collaboration Workshop to be held in Brussels in September 2011 [IoS2011].

The video can be seen on the SRT-15 project webpage (<http://www.srt-15.eu/>) and on the YouTube video sharing platform (<http://www.youtube.com/watch?v=M-Eo8lOGy3Y>)

References

- [CCRA] <http://www.enisa.europa.eu/act/rm/files/deliverables/cloud-computing-risk-assessment>
- [CGB10] Sunoh Choi, Gabriel Ghinita, and Elisa Bertino. A privacy-enhancing content-based publish/subscribe system using scalar product preserving transformations. In *Database and Expert Systems Applications*, volume 6261 of *Lecture Notes in Computer Science*, pages 368–384, 2010.
- [DGC] <http://debschallenge.event-processing.org/>
- [FIABud] <http://www.fi-budapest.eu/>
- [FIAGhent] <http://fi-ghent.fi-week.eu/>
- [HOLA] <http://www.holaproject.eu/>
- [IoS2011] http://ec.europa.eu/information_society/events/ssai/ios2011/index_en.htm
- [IoS2010] http://ec.europa.eu/information_society/events/ssai/ios/index_en.htm
- [PLAY] <http://www.play-project.eu/>
- [RM] <http://rapid-i.com/content/view/181/190/>
- [RTSS09] Thomas Ristenpart, Eran Tromer, Hovav Shacham, and Stefan Savage. Hey, you, get off of my cloud: exploring information leakage in third-party compute clouds. In *Proceedings of the 16th ACM conference on Computer and communications security, CCS '09*, pages 199–212, New York, NY, USA, 2009. ACM.
- [SM] <http://www.streamproject.eu/>
- [SMC] <http://smc.sourceforge.net/>

- [STREAM] <http://www.streamproject.eu/>
- [RDJF11] Romano L., DeMari D., Jerzak Z. and Fetzer C. A Novel Approach To QoS Monitoring In The Cloud. *1st International Conference on Data Compression, Communication and Processing*. Palinuro, Italy, 2011.
- [CCCDR11] Cicotti G., Coppolino L., Cristaldi R., D'Antonio S. and Romano L. QoS Monitoring in a Cloud Services Environment: The SRT-15 Approach. *Workshop on Cloud Computing: Projects and Initiatives*. Bordeaux, France, 2011.
- [TKG11] *Smart metering under EU Data Protection Law. International Data Privacy Law*. Trieb, Rainer Knyrim and Gerald. Oxford : First published online March, 1, 2011 doi:10.1093/idpl/ipr004, 2011.