SRT-15

Intelligence Push in the Enterprise Realm

RESEARCH NEWSLETTER

FALL 2012

IN THIS ISSUE

Zbigniew Jersak, Project Coordinator
SAP Research

EDITORIAL

A word from the project coordinator

The SRT-15 research project is achieving full maturity, as it is about to enter its third year of execution. During the last twelve months we have witnessed a series of achievements, which are highlighted in this newsletter.

One of the most important events that were co-organized by the SRT-15 project consortium was 6th ACM International Conference on Distributed Event Based Systems in Berlin, Germany. Prof. Christof Fetzer (TU Dresden) served as a Scientific Program co-chair while Zbigniew Jerzak (SAP AG) served as the DEBS 2012 Grand Challenge co-chair. The participation in the conference organisation strengthened also the ties of the SRT-15 research project with other EC funded research projects, most notably PLAY and KAP. Moreover, in the last twelve months the cooperation between consortium partners has significantly strengthened culminating with a donation of a 60 servers cluster from Yahoo! Iberia to the University of Neuchâtel.

The maturity of the SRT-15 research project is further strengthened by the availability of its components for public use (StreamMine3G, the PASC library). Finally, successful dissemination of the results is witnessed by high quality scientific publications (6th ACM International Conference on Distributed Event Based Systems, 15th International Conference on Extending Database Technology, 2012 USENIX Annual Technical Conference) as well as dissemination in the industrial community (SAP TechEd 2012 Las Vegas, Yahoo! Science Week 2012).

I would like to encourage you to browse through this newsletter and learn more about the relevance of the results of the SRT-15 research project to both industry and academia.

SRT-15 technologies

Tolerating data corruptions

One key dependability goal in the project was to harden key elements of the SRT-15 platform against random data corruptions generated by unreliable hardware. Most large-scale data processing platforms are designed to tolerate crashes; data corruptions can result in unexpected behaviour that violates the crash model and defeats existing fault-tolerance mechanisms. One infamous example is the 8-hours outage of the Amazon S3 service in July 2008: data corruption occurring in the state of a single server caused a system-wide outage due to error propagation [1]. Amazon reported that they use MD5 extensively but not for that particular part of the state. Similar failures occur in most companies handling large volumes of data, but are often kept confidential.

Distributed systems are typically hardened using ad-hoc error detection checks, a process that is time-consuming, requires application logic, and can easily overlook critical failure scenarios, as witnessed by the Amazon case. We developed a more principled and automated approach, which is based on a new fault model called ASC, for Arbitrary State Corruptions. ASC models the way hardware data corruptions impact the behaviour of processes of a distributed system. We also developed PASC, a hardening library that provably prevents error propagation in presence of ASC faults, ruling out the possibility of outages like the Amazon case. It has three key features: it is *transparent*, because it executes the necessary checks in the background without need of user intervention, *local*, because it does not require replication over multiple physical servers, and *untrusted*, because it guarantees error isolation even in presence of faults corrupting the state of the library itself.

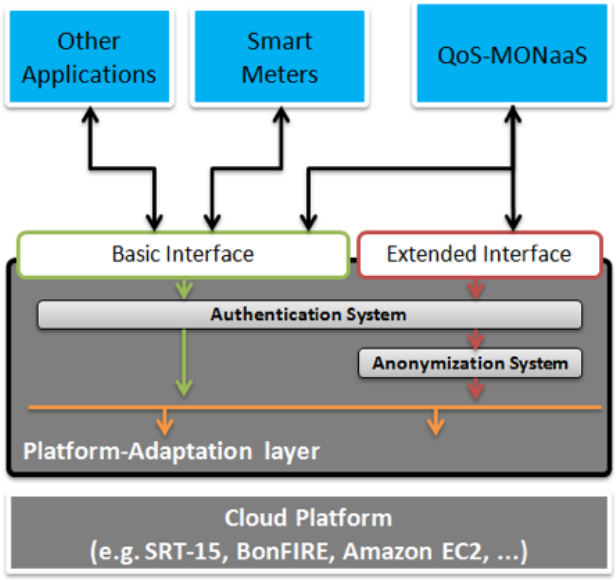

SRT-15
PILOT APPLICATION

A new version of the QoS-MONaaS pilot application is available and fully integrated with the current version of the SRT-15 platform. The application provides the ability to monitor Quality of Service (QoS) at the Business Process level.

We used PASC [2] to harden Zookeeper, the configuration service of the SRT-15 platform. We chose Zookeeper because it is a single point of failure of the system: a corruption of its state can make the whole system unavailable. The performance loss of hardening Zookeeper is acceptable: between 10% and 17% in terms of maximum throughput. We also conducted extensive fault injection campaigns to test the error isolation of PASC; we injected several thousands of random errors both in the code and in the state of Paxos, the underlying replication algorithm we used to implement Zookeeper, and found that PASC successfully prevents all error propagations.

[1] http://status.aws.amazon.com/s3-20080720.html
[2] https://github.com/yahoo/pasc

SRT-15 technologies

Privacy above all

One key aspect of the SRT-15 project is to enable the data to be sent through both private and public clouds, while ensuring the security of the content. Recent attacks in public clouds indicate that a major concern in untrusted domains is the enforcement of privacy. By routing data based on subscriptions evaluated on the content of publications, a content-based routing engine can expose critical information to unauthorized parties. In SRT-15, information leakage can be avoided by the means of privacy-preserving filtering, which is supported by several mechanisms for encrypted matching.

All existing approaches have in common a high performance overhead and the difficulty to use classical optimisations for content-based filtering such as per-attribute containment. In SRT-15, we use a novel mechanism that greatly reduces the cost of supporting privacy-preserving filtering based on encrypted matching operators. It is based on a pre-filtering stage and, when available, containment relationships between subscriptions, that greatly reduces the number of calls to the costly encrypted matching operator and significantly narrows the gap between cleartext and privacy-preserving filtering.

The pre-filtering operations leverage additional information attached to the subscriptions that encodes the values associated to the equality constraints of each subscription. Using truncation on the filters, our scheme gives very limited power to the attacker, not sufficient to run statistic-based attacks in practice and thus to break the privacy preservation of the non-optimised matching scheme.

Evaluations and experiments confirmed the ability of our mechanisms to significantly reduce the number of costly encrypted matching operations required to filter an incoming publication. Furthermore, our analysis shows that the additional data structures embedded in publications and subscriptions and used for pre-filtering have very limited impact on the effectiveness of privacy preservation.

This research has been published in the main technical track of DEBS 2012.

SRT-15 technologies

New version of QoS-MONaaS application

The ability of monitoring Quality of Service (QoS) at the Business Process level is key for a company's success, since assessing the actual quality of what service users are paying for has become a mission-critical business practice requirement, and it will be even more so in the future [1]. However, QoS monitoring entails a number of issues, since the ever increasing complexity of individual components, middleware, and interconnection infrastructures has impaired our ability to measure the Key Performance Indicators (KPIs) of the service which is to be monitored. Emerging development paradigms, and in particular cloud computing and Service Oriented Integration (SOI), together with the amazing increase of the scale, have made this task even more challenging in upcoming Future Internet (FI) scenarios, since individual business processes – whose internals are completely unknown – are being integrated to quickly implement and cheaply deploy semantically richer business processes. Evidence is demonstrating that QoS monitoring is and will be a very needed facility in the cloud computing scenario. In particular, in [2] authors emphasise that, due to the dynamic nature of cloud computing environments, continuous monitoring of Quality of

Service (QoS) attributes is needed to enforce Service Level Agreements (SLAs). In [3], authors state that "An independent tool for monitoring/validating performance of a heterogeneous set of applications" is one of the capabilities that are most needed in the cloud. Many companies are thus investing in the development of products that will provide QoS monitoring to enterprise services running in the cloud (including big players, e.g. IBM [4]). Epsilon, an Italian SME with strong commitment to research (10%+ of revenues invested in R&D), has developed — within the context of the SRT-15 project — an application that provides dependable QoS monitoring to enterprise services running in the cloud. The application is provided as a service itself. Thus, it is called QoS-MONaaS: Quality of Service MONitoring as a Service. QoS-MONaaS is being tested and validated with respect to a substantial case study of an Internet of Things (IoT) application developed by SAP within the context of the SRT-15 project. The case study, which is called Smart Meters, implements remote monitoring of power consumption in a Smart Grid environment.

QoS-MONaaS design was based on a clean-cut architecture, which clearly identifies (see figure on page 2):

- The interface that QoS-MONaaS shares with other applications (referred to as "Basic Interface"). This is the interface that all applications use to request platform services.
- The interface that QoS-MONaaS uses to gather information which is specifically needed for the purpose of QoS monitoring. This information is not provided to other applications. Thus, this interface is referred to as "Extended Interface".
- Two key services, namely Authentication and Anonymization, that QoS-MONaaS needs from the underlying platform (these are referred to as the "platform-dependent layer").

A stable version of QoS-MONaaS (QoS-MONitoring as a Service) pilot application, which incorporates Java Message Service (JMS) [5] support, is available and fully integrated with the current version of the SRT-15 platform.

A proposal for porting the QoS-MONaaS pilot on top of the BonFIRE testbed has been submitted to the 2nd BonFIRE Open Call [6]. Although it was not funded, the proposal has received a positive technical evaluation, particularly with respect to the architecture (the Evaluation Summary Report reads: "The plan to port the QoS-MONaaS monitoring solution from the SRT-15 project to the BonFIRE test-bed is sound").

[1] http://www.techout.com/resources/Why-Application-Monitoring-is-Critical-for-a-Companys-Online-Success.php

[2] P. Patel, A. Ranabahu, A. Sheth "Service Level Agreement in Cloud Computing", UKPEW 2009

[3] "Four Keys for Monitoring Cloud Services", White Paper from ManageEngine

[4] "Service_Level Agreement Monitoring with IBM Cognos Now!" , http://download.boulder.ibm.com/ibmdl/pub/software/data/sw-library/cognos/pdfs/factsheets/fs_sla_monitoring_with_ibm_cognos_now.pdf

[5] JMS, http://www.oracle.com/technetwork/java/index-jsp-142945.html

[6] BonFIRE Open Call, http://www.bonfire-project.eu/involved/open-call-info

SRT-15 partnership

Server donation

They have sent millions of message for two years, but they are now used for research and teaching at the University of Neuchâtel. Offered by Yahoo! Iberia, partner of the University of Neuchâtel in the SRT-15 project, these 60 high performance servers will be used for the development and large-scale evaluation of software for the SRT-15 project, as well as for teaching purpose, in the context of the new "Swiss Joint Master of Computer Science". These machines, which form a cluster, a representative infrastructure used in the industry, are a valuable tool for all researchers involved in the European project SRT-15. For additional details, please check the press release on the Web site of the SRT-15 project.

SRT-15 activities

Collaboration and dissemination

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.

The work on tolerating data corruptions has been published at the 2012 USENIX Annual Technical Conference in Boston. The results were presented at the Yahoo! Summer Science Week and in invited talks at HP Labs Palo Alto and Microsoft Research Silicon Valley.

The PASC library has also been open sourced and it is publicly available from https://github.com/yahoo/pasc.

UniNE issues a publication on efficient privacy filtering support in the CBR engine at DEBS 2012, and a publication on large-scale CBR engine scalability models at NCA 2012.

UniNE intends to use part of the technologies, know-how and methods developed in the context of SRT-15, in a new FP7 project on large-scale elastic architectures for data-as-a-service (LEADS). The project, to start in October 2012, will propose mechanisms for storing and searching private data along with public data in a federation of clouds.

SRT-15 contributed to the SPRERS (Strengthening the Participation of Romania at European R&D in Software Services) "2nd Training on Software Services — Cloud computing" event with a demo-based training session, which was held at West University in Timisoara, Romania on November 11-14 2011. EPSILON representative Luigi Sgaglione illustrated QoS-MONaaS main features to the trainees, and gave them the possibility to actively interact with the pilot. A video of the session is available from http://youtu.be/JKD6BXi0EN4.

SRT-15 HIGHLIGHT

Project member Christof Fetzer (TUD) acted as PC chair for DEBS 2012.

Christof Fetzer, DEBS 2012 PC Chair TU Dresden

DEBS 2012

The ACM International Conference on Distributed Event Based Systems (DEBS 2012) is the premier conference in distributed event processing. DEBS was this year organized at Freie Universitaet Berlin, July 16-20th, 2012, Berlin Germany.

SRT-15 member Christof Fetzer from TU Dresden was the Scientific Program Co-Chair of DEBS2012. The acceptance rate of DEBS 2012 was very low (21.5%) resulting in a high quality program. The technical program consisted of 17 full papers, 6 short papers, 6 industrial papers, 2 industrial experience reports, 6 demo and 5 poster papers, and 6 tutorials. These were complemented by 3 keynote talks provided by: Balan Sethu Raman (Microsoft), Martin Odersky (EPFL), and David Maier (Portland State University).

To place emphasis on the practical use of event-based technologies in distributed environments, the Grand Challenge was this year co-organised by the SRT-15 coordinator Zbigniew Jerzak from SAP Research. It was intended as a showcase of beneficial solutions for industry and commerce based on event-based technologies. The Grand Challenge was a great success and will be next year again co-organized by Zbigniew Jerzak.

SRT-15 members Raphaël Barazzutti, Pascal Felber, Hugues Mercier, Emanuel Onica and Etienne Rivière published their paper *Thrifty Privacy: Efficient Support for Privacy-Preserving Publish/Subscribe* in DEBS 2012.

DEBS 2012 Grand Challenge

The DEBS 2012 Grand Challenge was organized by the SRT-15 (grant number 257843, http://srt-15.eu) and PLAY (grant number 258659, http://www.play-project.eu) projects. The data and queries used within the DEBS 2012 Grand Challenge have been provided by the KAP project, grant number 260111, http://kap-project.eu. The DEBS 2012 Grand Challenge was hosted as a part of the 6th ACM International Conference on Distributed Event Based Systems, which took place in July in Berlin. The goal of the DEBS Grand Challenge series is to provide a common ground and evaluation criteria for a competition aimed at both research and industrial event-based systems. The goal of the DEBS Grand Challenge participants is to implement a solution to a specific problem provided by the DEBS Grand Challenge organizers.

The DEBS 2012 Grand Challenge focused on a use case that has been developed (together with the KAP project) based on the problems faced in a high-tech manufacturing industry. The DEBS 2012 Grand Challenge problem required a continuous monitoring of the high-tech manufacturing equipment, based on the data gathered by sensors embedded within the manufacturing equipment. The goal of the monitoring was to detect and record deviations from the predefined (good) system behaviour. The rationale behind the use of event based systems for monitoring of manufacturing equipment was the continuous nature of sensor data and a need for a low latency detection of possible violations.

The 2012 DEBS grand challenge had a total of 12 submissions out of which four have been accepted (after a peer review) to present during the 2012 DEBS Conference. All four accepted solutions have been included in the conference proceedings. The 2012 Grand Challenge issued awards in two categories: an overall Grand Challenge Winner Award (assigned by the Grand Challenge Program Committee and reviewers) as well as the Grand Challenge Audience Award (assigned based on votes cast by the Grand Challenge session audience members). The overall winner of the DEBS 2012 Grand Challenge has been the Homework system by Alexandros Koliousis and Joseph Sventek from the School of Computing Science at the University of Glasgow. The 2012 DEBS Grand Challenge Audience Award was given to Tilmann Rabl, Kaihwen Zhang, Mohammad Sadoghi, Navneet Kumar Pandey, Aakash Nigam, Chen Wang and Hans-Arno Jacobsen from the University of Toronto and University of Oslo for their submission titled "Solving Manufacturing Equipment Monitoring Through Efficient Complex Event Processing".

SRT-15 organised the DEBS 2012 Grand Challenge!