



SPEC ResearchSM Group Newsletter

CONTENTS OF THIS ISSUE

- 2** SPEC Research Group Officers
- 2** SPEC Research Working Groups
- 3** Foreword from Chairs and Editor
- 3** SPEC Research Group Mission Statement
- 3** SPEC Announcements
- 5** Report: Cloud Working Group
- 6** Report: IDS Benchmarking Working Group
- 6** Report: Big Data Working Group
- 7** Report: DevOps Performance Working Group
- 8** Three Distinctions of Alexandru Iosup
- 9** Winner of the SPEC Distinguished Dissertation Award 2015
- 10** Welcome to ICPE 2016 in Delft
- 10** ICPE 2017 Preliminary Announcement
- 10** LIMBO: a New Peer-reviewed Tool in SPEC Repository
- 11** Graphalytics: Benchmarking Platforms for Large-Scale Graph Processing
- 12** SPEC Symposium'16 – Energy Efficiency
- 13** InspectIT: Open Source APM Solution
- 14** diagnoseIT: A new APM Research Project
- 15** SPECtacular Awards 2015 and 2016

ICPE 2017 WILL BE HELD IN L'AQUILA, ITALY

Vittorio Cortellessa and Walter Binder, the General Chairs of the next ACM/SPEC International Conference on Performance Engineering (ICPE 2017), invite interesting high-quality submissions. The conference will take place in March 2017 in L'Aquila, Italy.

Read more on page 10

FOUR SPEC RESEARCH WORKING GROUPS REPORT ON THEIR PROGRESS

The SPEC Research Working Groups report on the progress, articles, benchmarks, and technical reports published in the year 2015. The Working Groups are always open for new members, feel invited to join us!

Read more on pages 5-7

YUFEI REN WINS THE SPEC DISTINGUISHED DISSERTATION AWARD

Among multiple high-quality dissertations that have been submitted to compete in the SPEC Distinguished Dissertation Award, the committee awarded Yufei Ren for his work on contemporary operating system virtualization techniques. Additionally, Cong Xu was distinguished with honorable mention.

Read more on page 9

NEW TOOLS IN THE SPEC REPOSITORY OF PEER-REVIEWED TOOLS

LIMBO and DynamicSpotter were accepted for the SPEC peer-reviewed tools repository. More tools are currently under review and further submissions are invited.

Read more on pages 3 and 10



CONTACT

Standard Performance Evaluation Corporation (SPEC)
7001 Heritage Village Plaza, Suite 225
Gainesville, VA 20155, USA

SPEC Research Group

Chair: Samuel Kounev (rgchair@spec.org)
Web: <http://research.spec.org>

SPEC RESEARCH GROUP OFFICERS

Chair:

Samuel Kounev, University of Würzburg, Germany,
rgchair@spec.org

Vice-chair:

André van Hoorn, University of Stuttgart, Germany

Secretary:

Klaus-Dieter Lange, HPE, USA

Release Manager:

Qais Noorshams, IBM, Germany

Steering Committee:

Samuel Kounev, University of Würzburg, Germany

Qais Noorshams, IBM, Germany

Tilmann Rabl, Bankmark UG, Germany

Petr Tůma, Charles University of Prague, Czech Republic

J. Nelson Amaral, University of Alberta, Canada

André van Hoorn, University of Stuttgart, Germany

Alexandru Iosup, TU Delft, The Netherlands

Lizy John, University of Texas at Austin, USA

Klaus-Dieter Lange, HPE, USA

Meikel Poess, Oracle Corporation, USA

Publicity Officer:

Christoph Heger, NovaTec Consulting GmbH, Germany

Newsletter Editor:

Piotr Rygielski, University of Würzburg, Germany

SPEC RESEARCH WORKING GROUPS

Cloud Working Group

Chair:

Alexandru Iosup, TU Delft, The Netherlands

Secretary:

Aleksandar Milenkoski, University of Würzburg, Germany

Release Manager:

Nikolas Herbst, University of Würzburg, Germany

<http://research.spec.org/working-groups/rg-cloud-working-group.html>

IDS Benchmarking Working Group

Chair:

Marco Vieira, University of Coimbra, Portugal

Vice-Chair:

Alberto Avritzer, Siemens Corporate Research, USA

Secretary:

Aleksandar Milenkoski, University of Würzburg, Germany

Release Manager:

Aleksandar Milenkoski, University of Würzburg, Germany

<http://research.spec.org/working-groups/ids-benchmarking-working-group.html>

Big Data Working Group

Chair:

Tilmann Rabl, Bankmark UG, Germany

Vice-Chair:

Todor Ivanov, Goethe University Frankfurt, Germany

Secretary:

Meikel Poess, Oracle Corporation, USA

Release Manager:

John Poelman, IBM, USA

<http://research.spec.org/working-groups/big-data-working-group.html>

DevOps Performance Working Group

Chair:

André van Hoorn, University of Stuttgart, Germany

Vice-Chair:

Felix Willnecker, fortiss GmbH, Germany

Secretary:

Robert Heinrich, KIT, Germany

Release Manager:

Alexander Wert, NovaTec Consulting GmbH, Germany

<http://research.spec.org/working-groups/devops-performance-working-group.html>

SPEC, the SPEC logo and the names SERT, SPEC SFS2014, SPEC, the SPEC logo and the names SERT, SPEC SFS, SPECjbb, SPECvirt_sc, Chauffeur WDK, and SPEC PTDaemon are trademarks of the Standard Performance Evaluation Corporation. The SPEC Research Logo and the name SPEC Research are service marks of SPEC. Additional company, product and service names mentioned herein may be the trademarks or service marks of their respective owners. Copyright © 1988-2016 Standard Performance Evaluation Corporation (SPEC). Reprinted with permission. All rights reserved.

WELCOME TO THE SPEC RESEARCH GROUP NEWSLETTER

We are delighted to present to you the next issue of the SPEC Research Group Newsletter. This regular publication provides information on latest developments, news, and announcements relevant to the benchmarking and quantitative system evaluation communities. Our newsletter is part of our mission to foster the exchange of knowledge and experiences between industry and academia in the field of quantitative system evaluation and analysis.

The past year has been very intense and exciting for the SPEC RG. To the major activities and milestones reached, we include: the organization of ICPE 2016, accepting two new tools for the SPEC repository of peer-reviewed tools, successful year for the four working groups, releasing two technical reports: (“New Cloud Metrics”, and “Performance-oriented DevOps: Research Agenda”), selecting the winner of the SPEC Distinguished Dissertation Award, and many more.

We have been actively working on preparation, planning and organization of the ICPE 2016. We hope that the vivid exchange of ideas during the upcoming ICPE 2016 will be a great motivation for the next year of scientific and engineering work.

We hope that you will enjoy reading the newsletter. We welcome and encourage your contributions for articles and suggestions for future coverage.

Samuel Kounev (University of Würzburg),
André van Hoorn (University of Stuttgart),
Piotr Rygliński (University of Würzburg)

SPEC RESEARCH GROUP MISSION STATEMENT

The SPEC Research Group (RG) is one of the four groups of the Standard Performance Evaluation Corporation (SPEC). Its mission is to promote innovative research in the area of quantitative system evaluation and analysis by serving as a platform for collaborative research efforts fostering the interaction between industry and academia in the field.

The scope of the group includes computer benchmarking, performance evaluation, and experimental system analysis considering both classical performance metrics such as response time, throughput, scalability and efficiency, as well as other non-functional system properties included under the term dependability, e.g., availability, reliability, and security. The conducted research efforts span the design of metrics for system

evaluation as well as the development of methodologies, techniques and tools for measurement, load testing, profiling, workload characterization, dependability and efficiency evaluation of computing systems.

Samuel Kounev (University of Würzburg)

SPEC ANNOUNCEMENTS

Zenodo repository for persisting ICPE 2016 submissions artifacts

January 20, 2016

Pursuing its mission to serve as a platform for collaborative research efforts, the SPEC Research Group (RG) has opened a new community on the Zenodo open digital repository. The purpose of the community is to archive and share research artifacts associated with RG activities, in particular the ACM/SPEC International Conference on Performance Engineering (ICPE). By providing a stable location for artifacts required to reproduce and reuse the research results, RG aims to broaden the practical impact of ICPE contributions and nurture an environment where detailed disclosure of research results is an accepted practice. At the time of writing, over a third of the authors presenting at ICPE 2016 has already uploaded the associated artifacts to the repository. More uploads are being prepared. The repository is available online under the provided link.

<http://zenodo.org/collection/user-spec-rg>

Petr Tůma (Charles University of Prague)

The SERT 1.1.1

January 18, 2016

The first quarter of 2016 sees the release of version 1.1.1 of the Server Efficiency Rating Tool (SERT) for the US EPA ENERGY STAR for Computer Servers program. This update adds support for the IBM Power8 processor family in little endian mode, validated with Ubuntu 14.04.03 and IBM J9 8.0 SR1.

It also includes a new revision of PTDaemon, v1.8.0. This adds support for the Yokogawa WT310E & 333E and the Hioki PW 3335, 3336 & 3337 power analyzers. All work in single channel AC mode, and also in multi-channel mode for those models that support more than one input channel.

As with every SERT update, the user guide has been thoroughly revised, based on user feedback and support questions received since the last release. This update is available in both English and Mandarin.

Mike Tricker (SPECpower Committee Secretary)

Article accepted for publication in the prestigious journal “ACM Computing Surveys”

July 28, 2015

A paper produced by the SPEC RG IDS Benchmarking Working Group has been accepted for publication in ACM Computing Surveys (CSUR). The paper surveys the state-of-the-art in the area of evaluation of intrusion detection systems.

CSUR publishes survey articles and tutorials related to computer science and computing. According to the ISI Journal Citation Reports, CSUR has among the highest impact factors of all computer science journals.

Aleksandar Milenkoski, Marco Vieira, Samuel Kounev, Alberto Avrtizer, and Bryan D. Payne. Evaluating Computer Intrusion Detection Systems: A Survey of Common Practices. ACM Computing Surveys, 2015, ACM.

www.research.spec.org

The 13th IEEE International Conference on Autonomic Computing (ICAC 2016)

September 29, 2015

The 13th IEEE International Conference on Autonomic Computing (ICAC 2016) will be held in Würzburg, Germany on July 19-22, 2016.

ICAC is the leading conference on autonomic computing, its foundations, principles, engineering, technologies, and applications. Nowadays, complex systems of all types, like large-scale data centers, cloud computing infrastructures, cyber-physical systems, the internet of things, self-organizing systems, organic computing systems, cognitive computing systems, or self-aware computing systems, are increasingly complex, involving many active, interconnected components requiring careful coordination. Being impossible for a human to manage such systems, the autonomic computing paradigm with its support for self-management capabilities becomes increasingly indispensable for the components of our IT world.

<http://icac2016.uni-wuerzburg.de/>

www.research.spec.org

SPEC SFS2014 benchmark received an update

November 19, 2015

This update fixes a possible sensitivity to systems that have inline compression enabled and incorporates all errata fixes since the initial release. Going forward,

all result submissions must be made using the SP1 version of the benchmark. Existing licensees of V1.0 are entitled to a free upgrade to SP1.

<https://spec.org/sfs2014/>

www.spec.org

DynamicSpotter was accepted for SPEC Tools Repository

January 21, 2016

A new tool was accepted for SPEC RG’s repository of peer-reviewed tools for quantitative system evaluation and analysis: DynamicSpotter. DynamicSpotter is a framework for measurement-based, automatic detection of software performance problems in Java-based enterprise software systems. DynamicSpotter combines the concepts of software performance anti-patterns with systematic experimentation.

More tools are currently under review and further submissions are encouraged.

<http://research.spec.org/tools/>

www.research.spec.org

SPECjbb2015 released

November 18, 2015

SPEC announces the release of SPECjbb2015, an enhanced Java server benchmark that now supports virtualized hosts, adds stricter performance validation, and specifies faster response times for a more accurate assessment of total capacity in a variety of real-world deployments.

<https://spec.org/jbb2015/>

www.spec.org

Book: Server Efficiency—Metrics for Computer Servers and Storage

June 19, 2015

Klaus-Dieter Lange, SPEC board member and SPECpower chair, together with Karl Huppler, a long-time supporting contributor to SPEC, are the authors of the recently released “Server Efficiency—Metrics for Computer Servers and Storage”, the 12th book in the ASHRAE Datacom Series. Huppler calls the book “a valuable guide for choosing energy-efficiency measures for computer servers and associated storage.” It can be purchased on the ASHRAE website.

<http://www.techstreet.com/ashrae/products/1894771>

www.spec.org

REPORT: CLOUD WORKING GROUP

January 18, 2016

The SPEC RG Cloud Group has continued, over the course of 2015, to pursue its long-term mission of taking a broad approach, relevant for both academia and industry, to cloud benchmarking, quantitative evaluation, and experimental analysis. Our group focuses on novel cloud properties such as elasticity, performance isolation, dependability, and other non-functional system properties, in addition to classical performance-related metrics such as response time, throughput, scalability, and efficiency. Among the keywords most discussed in the group, over 2015, are cloud metrics, elasticity, auto-scaling, service definition and configuration, and data as a service.

The scope of the group is ‘to develop new methodological elements for gaining deeper understanding not only of cloud performance, but also of cloud operation and behavior, through diverse quantitative evaluation tools, including benchmarks, metrics, and workload generators’. In 2015, through bi-weekly online meetings facilitated by WebEx and SPEC, we have advanced work on 4 main topics and a number of other sub-topics:

(1) Cloud Usage Patterns (CUPs): The goal of this activity is to define a formalism for expressing cloud usage patterns and scenarios. The joint authors, who represent seven organizations, propose [1] a simple yet expressive textual and visual formalism, which can be used by both general users and cloud experts. A key feature of the textual formalism is its conciseness; this goes in contrast to other formalisms that also focus on the executability of the specification. By expressing over ten patterns commonly seen in academic and industrial practice, the authors show that CUP is practical. An extended article, with new core concepts, was created during 2015 and is currently under submission.

(2) Cloud Metrics Survey and Design: This ongoing activity focuses on surveying existing cloud metrics and on the design of key missing metrics that allow the quantitative assessment and characterization of typical cloud usage scenarios. Among the key new metrics, the joint authors focus on various forms of elasticity [2] and risk-quantifying metrics [3]. A continuation of this work, and an extended set of metrics, have been developed by the group over 2015 and are being prepared for publication.

(3) BUNGEE: This ongoing activity focuses on the development and validation of a Java-based framework focusing on cloud elasticity, especially for IaaS cloud platforms and auto-scaling environments. The tool provides load and stress-testing functionality, and automates the analysis of the quality of the elastic behavior of the system under test through several elasticity

metrics. Currently, BUNGEE supports CloudStack and Amazon AWS based deployments. The group is currently extending BUNGEE to support more cloud environments and metrics.

(4) Benchmarking Auto-Scaling Techniques: This ongoing activity, which various members of the RG Cloud Group are just starting, is aiming to conduct a quantitative analysis and comparison of auto-scaling techniques in virtualized environments.

Current participants in the RG Cloud Group include the Delft University of Technology (Delft), the IBM T.J. Watson Research Center (USA), Lund University (Sweden), MITRE (USA), Oracle (USA), Salesforce.com (USA), SAP (Germany), Tata TCS (India), Umeå University (Sweden), and the University of Würzburg (Germany). Several of the members play leading roles in their institutions, them and others have won important distinctions over the course of 2015. For example, Dr. Samuel Kounev currently leads the Descartes Software Engineering Group of the University of Würzburg, and Dr. Alexandru Iosup has won the Netherlands Teacher of the Year 2015 award and has been selected to join the Young Academy of the Royal Netherlands Academy of Arts and Sciences. Besides the frequent participants, we have welcomed various invited speakers and guest participants. For 2016, we are actively seeking new participants and activities.

In 2015, the RG Cloud Group has been represented by its members at many international academic and industrial venues, and has helped organize various international events, such as HeteroPar and the flagship conference of SPEC, the ACM/SPEC International Conference on Performance Engineering (ICPE). Internally within SPEC, our group has contributed to the review and subsequent publication of several benchmarking tools. Concluding, 2015 was a full and successful year for the RG Cloud Group. We are looking forward to an even more successful 2016!

Acknowledgments: part of this work has been submitted by the group as a poster to the ICPE2016 conference, to inform our community about the opportunity of collaborating with the RG Cloud Group.

[1] A. Milenkoski, A. Iosup, S. Kounev, K. Sachs, P. Rygielski, J. Ding, W. Cirne, and F. Rosenberg. Cloud usage patterns: A formalism for description of cloud usage scenarios. CoRR, abs/1410.1159, 2014.

[2] N. R. Herbst, S. Kounev, and R. H. Reussner. Elasticity in cloud computing: What it is, and what it is not. ICAC’13, 2013, pages 23-27, 2013.

[3] V. van Beek, J. Donkervliet, T. Hegeman, S. Hugtenburg, and A. Iosup. Self-expressive management of business-critical workloads in virtualized datacenters. IEEE Computer, 48(7):46-54, 2015.

[4] N. R. Herbst, S. Kounev, A. Weber, and H. Groenda. BUNGEE: an elasticity benchmark for self-adaptive IaaS cloud environments. In SEAMS 2015, 2015, pages 46-56, 2015.

<http://research.spec.org/working-groups/rg-cloud-working-group.html>

Alexandru Iosup (TU Delft)

REPORT: IDS BENCHMARKING WORKING GROUP

January 14, 2016

The SPEC RG IDS Benchmarking Working Group successfully concluded its agenda for 2015 and faces 2016 with a renewed commitment. In 2015, the Working Group further established its research agenda having several works published.

The SPEC RG IDS Benchmarking Working Group published a paper focusing on evaluation of IDSes in virtualized environments at RAID 2015 (the 18th International Symposium on Research in Attacks, Intrusions and Defenses) [1]. This symposium brings together leading researchers and practitioners from academia, government, and industry to discuss novel security problems, solutions and technologies related to intrusion detection, attacks and defenses. The paper presents HInjector, a tool for generating IDS evaluation workloads by injecting such attacks during regular operation of a virtualized environment. It demonstrates the application of the tool and shows its practical usefulness by evaluating a representative IDS designed to operate in virtualized environments. The virtualized environment of the industry-standard benchmark SPECvirt_sc2013 is used as a testbed, whose drivers generate workloads representative of workloads seen in production environments. This work enables for the first time the injection of attacks in virtualized environments for the purpose of generating representative IDS evaluation workloads.

In addition, a paper produced by the SPEC RG IDS Benchmarking Working Group was published by ACM Computing Surveys (CSUR) [2]. The paper surveys the state-of-the-art in the area of evaluation of intrusion detection systems. CSUR publishes survey articles and tutorials related to computer science and computing. According to the ISI Journal Citation Reports, CSUR has among the highest impact factors of all computer science journals.

[1] Aleksandar Milenkoski, Bryan D. Payne, Nuno Antunes, Marco Vieira, Samuel Kounev, Alberto Avritzer, and Matthias Luft. Evaluation of Intrusion Detection Systems in Virtualized Environments Using Attack Injection. In The 18th International Symposium on Research in Attacks, Intrusions, and Defenses (RAID 2015), Kyoto, Japan, Springer. November 2015, Acceptance Rate: 23%.

[2] Aleksandar Milenkoski, Marco Vieira, Samuel Kounev, Alberto Avritzer, and Bryan D. Payne. Evaluating Computer Intrusion Detection Systems: A Survey of Common Practices. ACM Computing Surveys, 48(1):12:1-12:41, September 2015, ACM, New York, NY, USA. 5-year Impact Factor (2014): 5.949.

<http://research.spec.org/working-groups/ids-benchmarking-working-group.html>

Aleksandar Milenkoski (University of Würzburg), Marco Vieira (University of Coimbra), Nuno Antunes (University of Coimbra), Alberto Avritzer (Siemens Corporate Research, USA)

REPORT: BIG DATA WORKING GROUP

January 14, 2016

Big data is a rapidly evolving discipline in computer science utilizing a diverse spectrum of technical platforms and serving a wide range of applications. This is because, with the combinations of large volumes of data, heterogeneous data formats, and the rapidly improving performance of both hardware and big data systems, it is hard to generalize architectural aspects that best suit all application requirements, making the investigation and standardization of such systems very difficult.

As big data systems are evolving, there is an inherent need to evaluate and quantify their performance with the ultimate goal of comparing these systems. Comparisons are desirable in different dimensions, such as software stack, hardware, use case, and tuning parameters. That is, one might want to compare a particular software stack on different hardware systems, a particular hardware setting on different software stacks, or one software stack on a particular hardware with different tunings.

With the rapid increase in big data solutions, both academia and industry alike are developing new benchmarks at a rapid pace. Driven by the “velocity of change” many performance benchmark developers “cut corners” by customizing their benchmarks too closely to the architectural characteristic of the system they want to benchmark, instead of abstracting its core performance attributes. These benchmarks become “island solutions” that only fit the systems they targeted in the first place. This approach works well if the goal is to compare the performance of a particular software stack on a particular hardware setting. However, this does not work well to compare the performance of different software stacks on the same hardware platforms or vice versa.

Employing traditional ways of specifying standardized benchmarks with predefined workloads is not trivial for big data systems, because of the combinations of large volumes of data, heterogeneous data formats, and velocity of changes in the processing technology used in big data solutions. As a consequence, many companies and research institutions are developing their own “island solutions” that only fit systems they target. It is a challenge for both industry and academia to keep track of the large number of emerging benchmarks.

Therefore, in 2015, the Big Data Working Group focused on surveying current big data benchmarking efforts and published a paper on this topic at the Transaction Processing Performance Council’s Technical Conference (TPC TC) [1]. Furthermore, the group co-organized the 6th Workshop on Big Data

Benchmarking [2] in Toronto and held its first face to face meeting there. As an outreach activity, a poster and lightning talk was presented at the 8th Extremely Large Database Conference [3].

The group holds weekly meetings, with internal calls and open, public calls alternating. For the open calls, presentations on big data benchmarking, big data systems, performance tuning, and related research are invited. These calls are well attended and received. Topics in 2015 were BigBench, IoTAbench, big data in Finance, SparkBench, Oracle Big Data SQL, Graphalytics, HiBench, Benchmarking Facets, and benchmarking on modern clusters. All presentations can be found on the working group's website [4]. In the internal calls, the group furthers the research on big data benchmarking. In 2016, besides extending the ongoing research activities, the group plans to co-organize the 8th Workshop on Big Data Benchmarking, which will be held in Barcelona, Spain hosted by the Barcelona Supercomputing Center.

[1] Todor Ivanov, Tilmann Rabl, Meikel Poess, Anna Queralt, John Poelman, and Nicolas Poggi. Big Data Benchmark Compendium. Seventh TPC Technology Conference on Performance Evaluation & Benchmarking (TPCTC 2015).

[2] Sixth Workshop on Big Data Benchmarking, Toronto, Canada (WBDB2015.ca). <http://clds.sdsc.edu/wbdb2015.ca>

[3] John Poelman, Meikel Poess, Tilmann Rabl, Anna Queralt, Todor Ivanov, and Chaitan Baru. SPEC RG Big Data Working Group: An Introduction. 8th Extremely Large Databases Conference (XLDB '15).

<https://research.spec.org/working-groups/big-data-working-group/meetings-and-presentations.html>

Tilmann Rabl (bankmark UG)

REPORT: DEVOPS PERFORMANCE WORKING GROUP

December 21, 2015

The DevOps Performance Working Group [1] is concerned with combining application performance management (APM) and software performance engineering (SPE) activities for business-critical software systems. The need for a better SPE/APM integration is driven by an increased interrelation of development and operation teams in corporate environments, including new software development and organization paradigms known as DevOps. The DevOps movement aims on merging the classical development and operations silos to end-to-end responsible teams. This paradigm changes requires increased collaboration and a high degree of automation in the development and deployment processes. The community has failed to systematically integrate performance management practices into the DevOps world, even though modern software

development paradigms provide great opportunities to do that. Only holistic performance management, integrating Dev and Ops throughout the whole life-cycle, leads to high-quality software.

The DevOps Performance Working Group fosters and facilitates research in combining APM and SPE activities, e.g., by experience sharing, agreement on definitions, specification of metrics, and dissemination of novel methods, techniques, and tools for quantitative evaluation. The group focuses on the challenges and opportunities of performance evaluations in DevOps environments. We aim to identify cross-community collaboration, and to set the path for long-lasting collaborations towards performance-aware DevOps.

As the major joint activity in 2015, the DevOps Working Group created and published the technical report SPEC-RG-2015-01 "Performance-oriented DevOps: A Research Agenda" [2]. The report surveys performance evaluation methods, techniques, and best-practices for all lifecycle phases of enterprise software applications and identifies research challenges for the upcoming DevOps era. This report acts as a research agenda for the next years in this narrow domain related to different communities including performance and software engineering, cloud computing, and big data.

In February 2015, the group met for its first face-to-face meeting, co-located with the RELATE Winter School at the University of Würzburg, Germany. The agenda of the one-day meeting comprised detailed discussions on the technical report, two invited talks, and discussions on future working group activities. 13 group members participated in the meeting.

Together with two EU projects, the working group initiated the International Workshop on Performance-Aware DevOps (QUDOS). QUDOS 2015 [3], technically cosponsored by the SPEC RG, was collocated with the renowned 10th Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on the Foundations of Software Engineering (ESEC/FSE) in Bergamo, Italy in September 2015. The initial workshop comprised 10 talks, 8 scientific papers, and around 25 registered participants. Due to the great interest in this new workshop a second edition of QUDOS is planned for fall 2016.

Several sub-projects among subsets of the members have been conducted. A comparison and integration of resource demand estimation (LibReDE) and measurement (PMWT) techniques resulted in a publication at the 12th European Workshop on Performance Engineering (EPEW). Papers and tools on the WESSBAS approach for specifying, extracting, and generating session-based workloads have been published. A collaborative research project diagnoseIT on "Expert-Guided Automatic Diagnosis of Performance

Problems in Enterprise Applications” has been started (read more on page 14).

A current major joint activity is the development of a reference architecture and a process model for performance-aware DevOps. The technical report and a survey on the industrial DevOps performance practices, concurrently conducted by the working group, will serve to identify the core design rationale. The survey aims on detecting current DevOps practices in context of performance evaluation to align the reference architecture and process to industry practice. The outcome of this project will help companies to transform from classical development and operations silos to a performance-aware DevOps process model.

In September 2016, the group will co-organize a GI-Dagstuhl seminar entitled “Software Performance Engineering in the DevOps World” [4]. The main goal of the GI-Dagstuhl Seminar is to bring together young researchers in the areas of software engineering, performance engineering, cloud computing, and big data to present their current research projects, to exchange experience and expertise, to discuss research challenges, and to develop ideas for future collaborations

The group meets on a biweekly basis to discuss current activities and new project ideas, including invited presentations. In 2015, 22 meetings were held, including talks about performance problem detection, performance regression testing, and joint projects of the group. For more information about the DevOps Performance Working Group, (including mission, activities, meetings, presentations, projects) visit our web page [1]. If you are interested in following the discussions or contributing actively, please get in touch with the working group chairs.

[1] SPEC Research Group, DevOps Performance Working Group. Website

<http://research.spec.org/working-groups/devops-performance-working-group/>

[2] A. Brunnert, A. van Hoorn, F. Willnecker, A. Danciu, W. Haselbring, C. Heger, N. Herbst, P. Jamshidi, R. Jung, J. von Kistowski, A. Koziolok, J. Kroß, S. Spinner, C. Vögele, J. Walter, A. Wert: Performance-oriented DevOps: A Research Agenda. Technical Report SPEC-RG-2015-01, SPEC Research Group-DevOps Performance Working Group, Standard Performance Evaluation Corporation (SPEC), 7001 Heritage Village Plaza Suite 225, Gainesville, VA20155, USA, August 2015.

<https://research.spec.org/index.php?id=1345>

[3] 1st International Workshop on Quality-aware DevOps (QUDOS 2015).

<http://qudos2015.fortiss.org/>

[4] GI-Dagstuhl Seminar 16394, Software Performance Engineering in the DevOps World. Website.

<http://www.dagstuhl.de/16394>

André van Hoorn (University of Stuttgart),
Felix Willnecker (fortiss GmbH)

THREE DISTINCTIONS IN 2015 FOR ALEXANDRU IOSUP, OF THE CLOUD WORKING GROUP

January 18, 2016.

Alexandru Iosup, Associate Professor at the Delft University of Technology, has received three important distinctions in 2015. He was named Netherlands Teacher of the Year [1] on January 17, 2016, has been selected as a new member of the Young Academy of the Royal Netherlands Academy of Arts and Sciences (KNAW) [2] on November 25, 2015, and announced as the winner of the Netherlands Award for ICT Research 2016 [3] on December 17, 2015.

The title of Netherlands Teacher of the Year is awarded by the Interstedelijk Studenten Overleg (Inter-City Student Consultation, or ISO) to a teacher who is known for his or her involvement, passion, and innovative teaching methods. This is according to the jury the case for Iosup and his slogan ‘every student counts’. He has achieved high passing rates, high achievement, and high involvement in courses for which performance evaluation and benchmarking are important, such as the 1st-year Bachelors’ course Computer Organization, and the M.Sc.-level course Cloud Computing.

Dr. Iosup uses gamification in his teaching in order to help students learn in ways that suit their personalities. Students can follow their own routes to learning, thus earning credits. They may choose amongst exploratory, social, or competitive elements. Some students are fond of free assignments (e.g., designing and programming games), while others prefer to work in small groups and yet others seek the challenge of breaking through to more difficult material, like reaching a higher level in a game. In order to keep the students further involved (or ‘on board’), they may complete the entry quiz once, and they may participate in discussions on the newest and most important research and industrial challenges at the beginning of every lecture. This motivates students to go further and distinguish themselves. By creating challenges, excitement, a quest for a particular goal and rewards, Iosup keeps his students exceptionally involved. Many students are so motivated that they wish to take extra classes and do extra work for the sole purpose of learning something new, without any other reward; for these students Iosup organizes extra lectures and half-day workshops. The jury and his students were quite enthusiastic about Iosup: ‘He knows that each student is unique, and he is able to use gamification to bring out the best in all students.’

In November 2015, Alexandru Iosup has been selected with nine other scientists from other domains to join the Young Academy, which is an independent

science platform of leading young researchers. Every year, the Young Academy selects ten talented new researchers to add to its ranks. Members are elected for five years.

Dr. Iosup has been recognized for being an international pioneer in the field of distributed systems. His research was considered essential to emerging fields such as cloud computing, to popular internet applications such as e-Science and online gaming, and to the future of large-scale and/or gamification-based teaching. Iosup is an advocate of open access and devotes a great deal of his time to supervising talented young researchers. He is involved in master classes on technology, in particular in his field of research, for Dutch engineers and students.

In December 2015, Dr. Alexandru Iosup has been announced as the winner of the Netherlands Award for ICT Research 2016. The prestigious Netherlands Award for ICT Research is presented annually by the Royal Dutch Organization of Sciences (Koninklijke Hollandsche Maatschappij der Wetenschappen, KHMW, founded 1752) to a scientific researcher, not older than 40 years old, who has conducted innovative research or is responsible for a scientific breakthrough in the field of ICT. The prize is an initiative of the Royal Netherlands Academy of Arts and Sciences via ICT Research Platform Nederland (IPN), in collaboration with the Netherlands Organisation for Scientific Research (NWO). On 23 March 2016, Dr. Iosup will receive his award at a special ceremony during the ICT.OPEN 2016 event.

Dr. Iosup plans to use the €50,000 prize money to strengthen his research group in order to launch new research into graph-processing systems, starting with benchmarking and as a collaboration between the SPEC RG Cloud Group and the Big Data Working Group. The jury praised Dr. Iosup for his academic achievements, his active collaboration with companies, and his efforts to bring ICT research to the attention of a wide audience.

The SPEC Research Group wishes to commend Alexandru for these distinctions.

[1] <http://www.tudelft.nl/en/current/latest-news/article/detail/alexandru-iosup-nederlands-docent-van-het-jaar-1/>

[2] <http://www.dejongeakademie.nl/en/nieuws/young-academy-selects-ten-new-members>

[3] <http://www.tudelft.nl/en/current/latest-news/article/detail/alexandru-iosup-ontvangt-nederlandse-prijs-voor-ict-onderzoek-2016/>

Members of SPEC RG Cloud Working Group

YUFEI REN FROM STONY BROOK UNIVERSITY, NEW YORK WINS THE 2015 SPEC DISTINGUISHED DISSERTATION AWARD

January 12, 2016

The SPEC Distinguished Dissertation Award is an annual award that aims to recognize outstanding doctoral dissertations within the scope of the SPEC Research Group in terms of scientific originality, scientific significance, practical relevance, impact, and presentation. This year, the award committee considered 15 excellent dissertation submissions from universities around the world: Pennsylvania State University, Virginia Tech, Karlsruhe Institute of Technology, Australian National University, University College Dublin, University of Lisbon, Colorado State University, University of Malaya.

The winning dissertation from Yufei Ren of New York's Stony Brook University, nominated by Professor Dantong Yu, is about scalable end-to-end data I/O over enterprise and datacenter networks. The committee particularly appreciates the insights presented in this dissertation for improving data-intensive applications in data centers from a holistic perspective. The practically oriented work across multiple layers of the networked storage stack coupled with the solid experiments drew praise from the committee members. The award is to be handed over at the ICPE 2016 Conference in the Netherlands in March 2016.

In addition to selecting the winner, the committee decided to recognize one dissertation with honorable mention. This is the dissertation of Cong Xu, nominated by Professor Yuan Xie of the Pennsylvania State University, for contributions to modeling, circuit design and microarchitecture of emerging resistive memory.

The award selection committee for 2015 chaired by Prof. Lizy Kurian John of the University of Texas at Austin, consisted of the following members:

- Mauricio Breternitz, AMD, USA,
- Yao-Min Chen, Oracle, USA
- Ramesh Radhakrishnan, Dell, USA
- Seetharami Seelam, IBM, USA
- Petr Tuma, Charles University, Czech Republic
- Zhibin Yu, China Shenzhen Institute of Advanced Technology, CAS, China

The SPEC Distinguished Dissertation Award was established in 2011 to recognize outstanding dissertations within the scope of the SPEC Research Group in terms of scientific originality, scientific significance, practical relevance, impact, and quality of the presentation.

Lizy Kurian John (University of Texas at Austin)

WELCOME TO ICPE 2016 IN DELFT

January 12, 2016

The 7th ACM/SPEC International Conference on Performance Engineering (ICPE) 2016 which takes place in March 12-16, 2016 in Delft in The Netherlands with all its different tracks attracted a number of high-quality submissions selected in intensive review processes by different track committees.

In the research track we had 57 submissions with 19 being accepted for presentation at the conference. In the Visions and Work-In-Progress track, 6 out of 15 contributions were selected. In the industry track additional 7 articles out of 17 made it into the final conference.

In addition to those scientific articles the programme also includes a poster and demo track, tutorials, a doctoral symposium, and several interesting workshops, for example, the 2nd Workshop on Challenges in Performance Methods for Software Development (WOSP-C 2016); 5th International Workshop on Large-Scale Testing (LT 2016); 2nd International Workshop on Performance Analysis of Big Data Systems (PABS 2016).

We are glad to see you in Delft in March 2016!

Steffen Becker (TU Chemnitz), Alexandru Iosup (TU Delft)

ICPE 2017 IN L'AQUILA– PRELIMINARY ANNOUNCEMENT

January 12, 2016

The ACM/SPEC International Conference on Performance Engineering (ICPE) provides a forum for the integration of theory and practice in the field of performance engineering. ICPE is an annual joint meeting that has grown out of the ACM Workshop on Software Performance (WOSP) and the SPEC International Performance Engineering Workshop (SIPEW). It brings together researchers and industry practitioners to share ideas, discuss challenges, and present results of both work-in-progress and state-of-the-art research on performance engineering of software and systems. ICPE 2017 will be held in L'Aquila (Italy), around the beginning of April.

L'Aquila is the capital city of the Abruzzi region in Italy, and it is located approximately 100 kilometers east from Rome, with which it is connected by a highway through the mountains. Laid out within medieval walls upon a hillside in the middle of a narrow valley, tall snow-capped mountains of the Gran Sasso massif flank the town. A maze of narrow streets, lined with

Baroque and Renaissance buildings and churches, open onto elegant piazzas.

Home of the University of L'Aquila, it is a lively college town. The Department of Computer Science and Engineering, and Mathematics (DISIM [1]) is very active in the international research context, and in particular the Software Engineering and Architecture Group includes members that have contributed to the birth and the growth of WOSP first and then ICPE conferences.

The host institution will be the University of L'Aquila, in collaboration with Gran Sasso Science Institute (GSSI, [2]). GSSI was recently born in L'Aquila thanks to OSCE funding, and today it represents a key perspective for the city and the territory.

The contact person for ICPE 2017 is Vittorio Cortellessa [3], who will be General Co-Chair along with Walter Binder [4] from University of Lugano. Anne Koziolok [5] from Karlsruhe Institute of Technology, and Evgenia Smirni [6] from College of William and Mary will be Program Co-Chairs.

[1] DISIM, <http://www.disim.univaq.it>

[2] GSSI, <http://www.gssi.infn.it>

[3] Vittorio Cortellessa, <http://www.di.univaq.it/cortelle/>

[4] Walter Binder, <http://www.inf.usi.ch/faculty/binder/>

[5] Anne Koziolok http://sdq.ipd.kit.edu/people/anne_koziolok/

[6] Evgenia Smirni, <http://www.cs.wm.edu/~esmirni/>


Vittorio Cortellessa (University of L'Aquila, Italy)

LIMBO: A TOOL FOR MODELLING, EXTRACTING AND GENERATING TIME-VARYING LOAD INTENSITY PROFILES

January 14, 2016

LIMBO is an Eclipse-based tool for the creation of load profiles with variable intensity over time both from scratch and from existing data [1]. Primarily, LIMBO's intended use is the description of load arrival behaviors in open workloads. Specifically, LIMBO can be employed for the creation of custom request or user arrival time-stamps or for the re-parameterization of existing traces.

LIMBO bases on the Descartes Load Intensity Model (DLIM) for the formalized description of its load intensity profiles. The DLIM formalism can be understood as a structure for piece-wise defined and combined mathematical functions that capture specific elements of a time-series like seasonal patterns, trends, bursts and noise.



LIMBO shows the decomposition of a load intensity profile.

LIMBO is capable of generating request or user arrival time stamps from DLIM instances.

In a next step, these generated time-stamps can be used for both open workload-based benchmarking or simulations. For example, TimestampTimer plug-in for JMeter already allows the precisely timed sending of requests to a system-under-test.

LIMBO offers a range of features for easy load intensity modeling and analysis, including, but not limited to, a visual decomposition of load intensity time-series into seasonal and trend parts, a simplified load intensity model as part of a model creation wizard, and an automated model instance extractor.

LIMBO's development has been started in late 2013 at the Karlsruhe Institute of Technology and is now maintained at the University of Würzburg.

Since LIMBO was accepted for the SPEC RG's tool repository in 2014, it has been adopted in research projects like EU FP7 projects Cactus and CloudScale. In addition, LIMBO has become a foundation for a Cloud Elasticity Benchmarking Framework named BUNGEE [2].

[1] Jóakim von Kistowski, Nikolas Roman Herbst, Daniel Zoller, Samuel Kounev, and Andreas Hotho. Modeling and Extracting Load Intensity Profiles. In Proceedings of the 10th International Symposium on Software Engineering for Adaptive and Self-Managing Systems (SEAMS 2015), Firenze, Italy, May 18-19, 2015.

[2] Nikolas Roman Herbst, Samuel Kounev, Andreas Weber, and Henning Groenda. BUNGEE: An Elasticity Benchmark for Self-Adaptive IaaS Cloud Environments. In Proceedings of the 10th International Symposium on Software Engineering for Adaptive and Self-Managing Systems (SEAMS 2015), Firenze, Italy, May 18-19, 2015

<http://descartes.tools/limbo>

Jóakim v. Kistowski and Nikolas Herbst (University of Würzburg)

GRAPHALYTICS: BENCHMARKING PLATFORMS FOR LARGE-SCALE GRAPH PROCESSING

January 18, 2016

Graphs model social networks, human knowledge, and other vital information for business, governance, and academic practice. Although both industry and academia are developing and tuning many graph-processing algorithms and platforms, the performance of graph-processing platforms has never been explored or compared in-depth. Moreover, graph processing exposes new bottlenecks in traditional HPC, distributed, multi-/many-core, and multi-threaded systems (for example, see the major differences in Top500 and Graph500 rankings).

Complementing other tools for benchmarking graph-processing systems, the Delft University of Technology, member of the RG Cloud Group, is leading a joint benchmarking activity that has already resulted in the creation of the Graphalytics benchmark and open-source tools [1,2]. Graphalytics focuses on new methods and tools for performance evaluation and optimization for graph processing platforms, and in particular on enabling its users to explore and to explain the performance dependency Platform-Algorithm-Dataset. It has already resulted in new knowledge and methodology [3]. One of the key results in 2015 has been the design and development of the Granula performance component for fine-grained performance data collection and analysis, which offers for big data processes unique features for the analysis of bottlenecks and resource usage.

Benefiting from a generous gift from SPEC member Oracle, Graphalytics is now mature software, and since

2015 is available to the public as a GitHub open-source software [2]. It has been tested by the Graphalytics development team, and by a multi-institution development team affiliated with LDBC [3], including Neo4j and SPEC members Oracle, Intel, and IBM (ongoing). It has also been used in various academic and industry projects, such as MSc thesis projects at the Delft University of Technology. In 2016, it aims to become a SPEC-endorsed tool.

In 2015, the Graphalytics team has presented the methods, tools, and core results also as tutorials, at ICPE 2015 in January (preliminary features), and at SCI15 in November (mature tools and optimization processes). The tutorials focused on real-world experiences with commonly used systems, from graph databases such as Virtuoso and Neo4j, to parallel GPU-based systems such as Totem and Medusa, to distributed graph-processing platforms such as Giraph and GraphX. The tutorials have included significant hands-on and Q&A components, but also in-depth analysis of system performance and bottlenecks.

We invite the SPEC Newsletter audience to try Graphalytics and contact our team.

Contact: A.losup@tudelft.nl

[1] Graphalytics Website: <http://graphalytics.ewi.tudelft.nl>

[2] Graphalytics Code: <https://github.com/tudelft-atlarge/graphalytics/>

[3] Mihai Capota, Tim Hegeman, Alexandru Iosup, Arnau Prat-Pérez, Orri Erling, Peter A. Boncz: Graphalytics: A Big Data Benchmark for Graph-Processing Platforms. GRADES@SIGMOD/PODS 2015: 7:1-7:6

Alexandru Iosup (TU Delft)

SPEC SYMPOSIUM'16–ENERGY EFFICIENCY: FREMONT, CA JANUARY 11TH, 2016

January 18, 2016

The inaugural SPEC's Symposium'16–Energy Efficiency was held in Fremont, California in January 2016. We were delighted to having brought together representatives from the U.S. Environmental Protection Agency, The Green Grid, ISO, ITIC, academic researchers, and SPEC benchmark developers in order to share ideas, discuss challenges, and present results of both work-in-progress and state-of-the-art research on server efficiency. It was with great honor to welcome the keynote speaker Jeremy Arnold (IBM) and listening to his speech "What is a Good Benchmark?". He discussed what makes a good benchmark, and how benchmarks should be selected and used based on the performance measurement goals of the user.

Sanjay Sharma (Intel) described the history and



Sanjay Sharma (SPECpower Committee Vice-Chair).

evolution of SPECpower_ssj2008, the industry's first server energy efficiency benchmark.

An overview of all SPEC's Energy Metrics were presented by David Schmidt (Hewlett Packard Enterprise)

Hansfried Block (Fujitsu) summarized the features of the Server Efficiency Rating Tool (SERT), together with the history of its evolution and the current use with the EPA ENERGY STAR program today.

John Clinger (ICF International/EPA) provided insight into the possible future evolution of the U.S. EPA's ENERGY STAR for Computer Servers program. He also provided ideas for how the use of the SERT might evolve, including the addition of metrics.

Jóakim von Kistowski described in his session, Research and SPEC Chauffeur, how he and his research colleagues at the University of Würzburg are using the Chauffeur WDK (derived from the SERT program) to rapidly develop new tools for measuring energy efficiency.

Xiaomeng Lu shared her insight, connections between industry and policymakers, into the challenges western organizations face when interacting with Asia Pacific governments and policy-makers on energy efficiency and benchmarking.

Henry Wong (Intel, ISO & TGG) described in his presentation ISO and Global Government Programs, how ISO can help to minimize the political aspects of

benchmark and tool sharing across world markets. He aligned the discussion with Xiaomeng's description of the Asia Pacific challenges.

The panel "Panel Discussion: The Future of Energy Efficiency of Server" with John Clinger, Henry Wong, Xiaomeng Lu, and Klaus-Dieter Lange discussed the premise: "The dawn of performance engineering is over". The discussion covered the tools and programs together with their ideas for future international collaboration.

Klaus and Xiaomeng delivered a SERT Demonstration on how the SERT is used to measure the energy efficiency of a server, using a server, a laptop, a power analyzer, and a temperature sensor.

The second presentation of the SPEC Symposium'16 takes place in Delft in the Netherlands. This event is co-located with the International Conference on Performance Engineering (ICPE) 2016, at the TU Delft in the University Aula. This symposium includes content derived from that presented in Fremont in January, along with EU-centric sessions and a panel discussion. It is hoped that industry and academic representatives who were unable to attend the US event will be present. Representatives of several Asian Government bodies and standards organizations have been invited. They are presenting content specific to their national objectives for international cooperation on energy efficiency standards and testing.

Mike Tricker (SPECpower Committee Secretary)

inspectIT THE OPEN SOURCE APM SOLUTION

January 5, 2016

Today's businesses have understood what impact the performance and scalability of their software applications has on revenue and total cost of ownership. These businesses use application performance management solutions to maintain the satisfying level of service quality, to detect problems early and to diagnose problems quickly to insure a stable revenue stream and reasonable total cost of ownership. Some businesses remain with an unsatisfied need to manage the performance of their applications. Because many businesses do not want to take the risk of a vendor lock-in that may dissolve or discontinue the product or cannot cope with the acquisition and annual maintenance costs. The public sector on the other hand may not be able to make a benefit of using the product due to budget restrictions or if the source code cannot be examined for safety or security reasons. Research groups are also often not able to make use of the

commercial solutions because of the high acquisition costs. In light of these observations, the matured tool inspectIT was made public to provide an open-source application performance management solution.

Following this mission, inspectIT is the first open-source tool of its kind. Software performance experts can monitor execution traces from applications under analysis and drill down into traces to isolate the root causes of performance problems. The rich graphical user interface pays attention to deliver great user experience designed to support the tasks of a performance expert's daily business in diagnosing problems. In the last decade, inspectIT has proven its reason for existence in many customer projects in industry. The recent collaboration with research groups shows the need for an independent performance analysis tool. Providing a standardized data interface enables researchers to collect data, export this data, share the data with other groups and import the data of others for analysis. This objective is already supported by the collaboration with researchers of different universities to standardize a non-proprietary data format for traces (currently referred to as the Common Trace API—CTA). The data format is used to export traces from inspectIT and to import traces from other tools of the community.

The intended users of inspectIT are performance experts and developers. By design, testing and troubleshooting scenarios (so called performance firefights) are the primary task which is intentionally supported by inspectIT. For this particular purpose, inspectIT provides a sophisticated graphical user interface that is tailored based on the experience of a decade to allow easy access to the collected performance data (e.g., traces, timings, CPU and memory utilization, etc.) and quick isolation of a problem's root cause. inspectIT is divided into three main components: (1) the agents, (2) the Central Measurement Repository, and (3) the graphical user interface. The agents are responsible for collecting the data. inspectIT currently provides an agent for Java-based applications. The Java agent is started with the system under analysis (SUA) and is responsible for instrumenting the system to collect the desired data. The agent is currently capable of collecting HTTP requests, execution traces, timings (e.g., response times, CPU times), thrown exceptions, SQL queries and system measures (e.g., CPU utilization, Memory utilization). One or more agents, depending on the concrete environment, send the collected data to the Central Measurement Repository which is responsible for managing the received data. The graphical user interface enables the performance expert to access the managed data and to execute performance analyses such as digging into traces to isolate the root cause of a problem, browse thrown exceptions, and view resource utilization. The user interface also shows the effective

instrumentation of SUA and allows to configure the instrumentation by associating namespaces, classes or interfaces of the SUA with the particular type of required data (e.g., timings, execution trace).

inspectIT is released under the GNU AGPL v3 license. The quality of inspectIT is a major objective that drives the integration of any change committed to the project's code base. The source code is managed on Github [1] together with ready to use Docker containers [2]. While the current release is inspectIT 1.6, further releases including bug fixes and new features are already planned in the roadmap for the upcoming major releases.

For more information about the inspectIT project and how to contribute to its development, visit our web page and get in touch with us.

[1] <https://github.com/inspectIT>

[2] <https://github.com/inspectit-docker>

Patrice Bouillet, Christoph Heger, Ivan Senic, Stefan Siegl, Alexander Wert (NovaTec Consulting GmbH)

NEW APM RESEARCH PROJECT: diagnoseIT

January 5, 2016

In the last decade, application performance management (APM) solutions have been developed supporting enterprises with monitoring capabilities and early detection of performance problems. Leading application APM solutions mostly support only alerting and visualization of performance-relevant measures. The configuration of the software instrumentation, the diagnosis of performance problems, and the isolation of the concrete root cause(s) often remain error-prone and frustrating manual tasks. To this day, these tasks are performed by costly and rare performance experts. In order to improve this situation, NovaTec Consulting GmbH and the University of Stuttgart (Reliable Software Systems Group) launched the collaborative research project diagnoseIT on “Expert-guided Automatic Diagnosis of Performance Problems in Enterprise Applications”. Hereby, the core idea is to formalize APM expert knowledge to automatically execute recurring APM tasks such as the configuration of a meaningful software instrumentation and the diagnosis of performance problems to isolate their root cause. By delegating the described tasks to diagnoseIT, experts do not have to deal with similar problems over and over again. Instead, the expert can focus on more challenging (and interesting) tasks.

The automated diagnosis that analyzes traces (individual or as a stream provided by monitoring capabilities of APM solutions) is designed as follows: possible symptoms of performance problems are provided as formalized expert knowledge—an extensible set of rules. When a symptom is detected in a trace, the root cause diagnosis is started without the need for human interaction. Rules that perform localization of the problem are applied first, followed by technology and/or domain-specific rules, which are used to semantify the isolated root cause. If the available data is not sufficient to perform an exhaustive diagnosis, diagnoseIT will request more data from the underlying APM tool, with respect to overall system performance.

diagnoseIT provides the analysis results to its user in the form of comprehensive reports that include qualitative (e.g., problem's location, type, and anti-pattern) and quantitative information (impact of the problem in numbers). A major goal of the report is to describe the problem to non-experts making it possible to provide different report types to individual roles.

The available APM solutions have a proprietary data format to persist and process monitoring data that makes it difficult and inconvenient, if not impossible, for third-party access. In order to enable community collaboration, data exchange, and tool-independent approaches in research, our goal is to make diagnoseIT independent of concrete APM solutions by accessing the monitoring data through the use of an independent API (currently referred to as the Common Trace API—CTA) and adapters for particular tool-dependent data formats. The specification of the CTA is an ongoing task motivated by serving as foundation for community collaboration on APM interoperability.

diagnoseIT started in March 2015 with a duration of two years. It is sponsored by the German Federal Ministry of Education and Research (BMBF), and supported by the SPEC RG DevOps Performance Working Group. For more information about the diagnoseIT project, as well as the CTA and how to contribute to its development, visit our web page and get in touch with us.

<http://diagnoseit.github.io/>

Christoph Heger, Alexander Wert, Stefan Siegl (NovaTec Consulting GmbH), André van Hoorn, Dušan Okanović (Reliable Software Systems Group, University of Stuttgart)



SPECtacular awards await the awarding ceremony in Freemont, 2016.
Photo: David Reiner (SPEC)

SPEC AWARDS IN YEAR 2015 AND 2016

January 18, 2016

During SPEC RG's 2015 Yearly Meeting in Austin (Texas, USA), RG members John Murphy, Klaus-Dieter Lange, Walter Binder and André van Hoorn were recognized by the SPEC President for their outstanding service to SPEC in 2014.

John Murphy (University College Dublin) was awarded with the SPEC 2014 Presidential Award for serving as General Chair of the ICPE 2014 recognizing his contributions to the success of the conference. Klaus Dieter Lange (HPE) and Walter Binder (University of Lugano, Switzerland) were recognized with SPECtacular Award for serving as co-general Chair for ICPE 2014. André van Hoorn (University of Stuttgart) received SPEC Award for project leadership in the DevOps Performance Working Group.

The year 2015 was also rich in events and activities that required outstanding work and commitment of Research Group members. During the

annual meeting in Freemont, California, USA (January 2016), two Research Group members were awarded with SPECtacular Awards: Aleksandar Milenkoski (University of Würzburg) for his excellence and commitment in IDS Benchmarking Working Group, and Kai Sachs (SAP) for serving as Program Committee Chair for ICPE 2015 held in Austin, Texas.

We congratulate all the winners and thank for their excellent support for SPEC Research Group.

Piotr Rygielski (University of Würzburg),
Alan MacKay (IBM) Walter Bays (Oracle)



André van Hoorn receives the SPECtacular award from SPEC President Walter Bays. Austin 2015. Photo: David Reiner (SPEC)



Klaus-Dieter Lange was SPECtacular and receives the award from SPEC President Walter Bays. Austin 2015.
Photo: David Reiner (SPEC)



Aleksandar Milenkoski (University of Würzburg) and Kai Sachs (SAP, Germany) were SPECtacular. Jóakim von Kistowski receives the awards in the name of the awarded. Freemont 2016.
Photo: David Reiner (SPEC).

