



RH RQ

Bringing great research ideas
into open source communities

Michael Zink

*On shared cloud
computing resources
making research more
accessible and powerful*



Optimizing Kubernetes

Ops is the new code

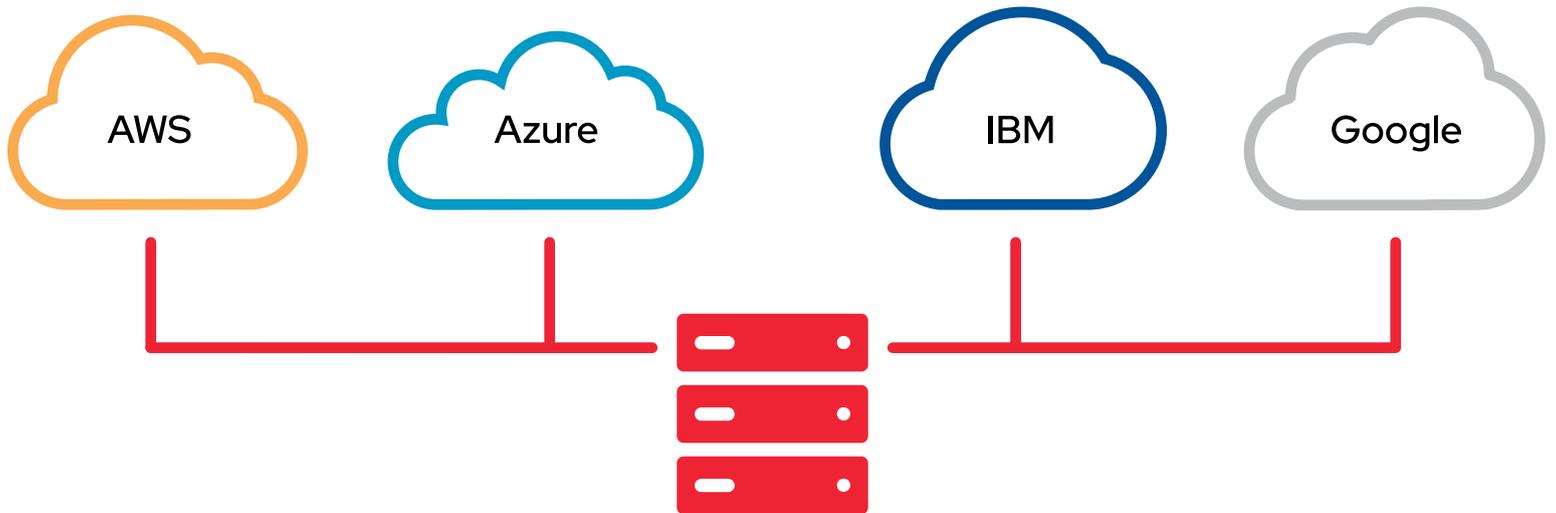
Where will we find
the data scientists?



Red Hat
Research Quarterly

Volume 3:4 | February 2022 | ISSN 2691-5278

Clouds that compete



should still connect.

| Our code is open_



Red Hat

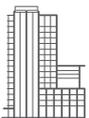
redhat.com/ourcode

Table of Contents



Departments

- 04** From the director
- 05** News: Building a university OSPO
- 07** News: Papers by partners in Europe achieving recognition
- 08** News: Call for code
- 12** News: Red Hat Collaboratory Research Incubation Award winners
- 16** News: Open source researchers win 2021 innovation awards
- 40** PhD support in the Czech Republic
- 42** Research project updates



ABOUT RED HAT Red Hat is the world's leading provider of open source software solutions, using a community-powered approach to provide reliable and high-performing cloud, Linux®, middleware, storage, and virtualization technologies. Red Hat also offers award-winning support, training, and consulting services. As a connective hub in a global network of enterprises, partners, and open source communities, Red Hat helps create relevant, innovative technologies that liberate resources for growth and prepare customers for the future of IT.

NORTH AMERICA
1 888 REDHAT1

EUROPE, MIDDLE EAST,
AND AFRICA
00800 7334 2835
europe@redhat.com

ASIA PACIFIC
+65 6490 4200
apac@redhat.com

LATIN AMERICA
+54 11 4329 7300
info-latam@redhat.com

Features

- 18** Clouds on the horizon: an interview with Michael Zink
- 26** Optimizing Kubernetes service selection
- 32** Ops is the new code
- 36** Where will we find the data scientists?



facebook.com/redhatinc
@redhatnews
linkedin.com/company/red-hat

From the Director



About the Author

Hugh Brock is the Research Director for Red Hat, coordinating Red Hat research and collaboration with universities, governments, and industry worldwide. A Red Hatter since 2002, Hugh brings intimate knowledge of the complex relationship between upstream projects and shippable products to the task of finding research to bring into the open source world.

Operating without borders

by *Hugh Brock*

If you take a moment to look at open source – the process, the language and ways of thinking involved, the legal framework – you will quickly realize that it is quite similar to the model for the free advancement of knowledge and thinking that the world’s universities have developed over the last thousand years. At the core of both models is the notion that the fastest way to advance our knowledge of the world is to share learning and results openly and acknowledge that there can be truth in a paper or a finding or a piece of code that is provable and not subject to political criticism or the whims of fashion.

Despite this similarity, I spend a lot of time, in these pages and in my continuing conversations with researchers and teachers, talking about how important it is to bring open source to universities and to the learning process in general. Why is this necessary? I think it has less to do with the openness of the research process than it does with the institutions that surround it. Because they are institutions, universities create hierarchies with walls around them. In doing so, they build barriers to participation in the research process.

Open source, at least when it’s done well, is engaged in leaping over those barriers and allowing anyone to contribute who cares to make the effort. From this stems the core mission of Red Hat Research. By working with university researchers to connect their research findings with the software communities we participate in, we can both ensure the continued relevance of those findings and engage people outside the university walls in the research process, as they should be.

This issue features several examples of progress in connecting researchers with open source methods and resources. First, we have just announced the first round of research grant awards from the \$4-million-per-year Red Hat Collaboratory grant to Boston University. This is one of the largest annual university grants in the United States targeted specifically at research with an open source component, and I’m really excited to be watching the results from it unfold. Each project we fund has at least one dedicated Red Hat engineer involved to help ensure the results make their way into open source. See the descriptions in the News section for details.

Another of the major open source-connected research institutions in the United States is the Center for Research in Open Source Software (CROSS), funded by retired Red Hatter Sage Weil at the University of California Santa Cruz. Stephanie Lieggi's piece on CROSS's Open Source Program Office is another great study in how to pull open source and research together.

A big part of open source innovation, particularly at Red Hat, is the computing infrastructure researchers and corporations use to do their work. Professor Michael Zink, our interviewee this issue, has been working with open source infrastructure for decades and has a lot to say about the various NSF- and university-funded efforts in this space. As we transition – hopefully – to operations that are open the way code is now open, Zink's experience and that of others like him will be critical to making community-managed infrastructure possible.

Speaking of open operations, Gordon Haff's piece on Operate First describes what we are doing at Red Hat Research, Boston University, and the Mass Open Cloud to make operating with complete transparency possible. Open operations must be the future of open source, particularly as AI becomes more important to operations; Haff explains why this is and what we are doing to help. 

Building a university OSPO: Bolstering academic research through open source

by Stephanie Lieggi

In industry, companies establish open source program offices (OSPOs) to mitigate potential legal risks, improve engineering practices, and enable financial benefits. OSPOs are also proving to be important in research university settings, although—not surprisingly—for different reasons. The primary aim of an academic OSPO is to advance the institution's goals, which typically include education, research, and serving the wider community.

At the University of California, Santa Cruz, the [Center for Research in Open Source Software \(CROSS\)](#) increases the impact of university research by strengthening relationships in open source ecosystems. Building on the success of CROSS, UC Santa Cruz is now establishing the first OSPO in the University of California system and taking huge strides towards establishing a system-wide organization, which would potentially be the largest academic-based OSPO to date.

A SHORT HISTORY OF CROSS
CROSS started after UC Santa Cruz alumnus Sage Weil sold his startup, built

around open source storage project Ceph, to Red Hat. Weil, who developed Ceph while a PhD student at UC Santa Cruz, wanted other students to have the same opportunity he did. In 2015, Weil gifted \$2 million to UC Santa Cruz, and specifically to Ceph co-creator and UC Santa Cruz professor Carlos Maltzahn, to create a center that would bridge the gap between student work and successful open source projects. With Weil's gift and three founding industry members—Toshiba (now Kioxia), SK Hynix, and Micron—Maltzahn launched CROSS as a research center under UC Santa Cruz's Baskin School of Engineering. The Center's purpose is to support innovative research and incubator projects that leverage open source techniques and communities in a university setting.

Over the last six years, CROSS has shown that successful open source projects are a powerful catalyst for industry engagement that also significantly increases the impact of university research. The CROSS team, led by Maltzahn and Assistant Director Stephanie Lieggi, has raised \$2.6 million



About the Author

Stephanie Lieggi

is the Assistant Director at CROSS, supporting the work of open source projects and enabling a sustainable contributor base through the establishment of hands-on mentorship programs. She promotes the use of open source in academic settings as well as increasing diversity and inclusion in open source ecosystems. Stephanie was a Senior Research Associate and Adjunct Professor at the Center for Nonproliferation Studies at the Middlebury Institute of International Studies at Monterey.

in membership fees from eight corporations to fund ongoing work, including nineteen graduate fellows' projects and five incubator projects led by postgraduate fellows. Most recently, CROSS's flagship incubator project, [SkyhookDM](#), was successfully merged into the Apache Arrow mainline and included as part of the Arrow 7.0.0 release. Other incubator projects have successfully concluded after creating sustainable and healthy communities for their continued development.

The graduate and postdoctoral students supported by CROSS have mentored student contributors and provided exciting and relevant projects for open source newcomers. CROSS fellows' student mentorship and industry engagement have helped with seeding their projects' contributor communities and teaching a wider pool of students how to engage in open source productively. To fully capture this dynamic, CROSS established its Open Source Research Experience (OSRE) program in summer 2020. Based on the Google Summer of Code (GSoC) model, the OSRE matches students with mentors for short-term projects that bolster the work of UC-based open source projects while providing students hands-on experience guided by expert mentors. OSRE also matches mentors and students with interested industry sponsors who support their work through an open source sponsorship agreement.

OSPO: THE NEXT STEPS

Based on the experience of the last six years, the CROSS team is now moving past the initial objectives of the Center. CROSS is also strategizing to bridge other significant gaps, including insufficient opportunities for student training in open source, lack of industry engagement with academic researchers in open source ecosystems, and absence of an effective method for assigning value to open source in university settings.

In response, the team at UC Santa Cruz, inspired by steps taken at other schools like John Hopkins Univeristy and the Rochester Institute of Technology, as well as the formation of the OSPO++ group in early 2020, began laying the foundations of an open source program office. In November 2021, they were awarded a two-year grant from the Sloan Foundation—the same foundation that supported OSPOs at JHU and RIT—to create a pilot project furthering this effort. The funding underwrites a newly imagined open source incubator fellows program as well as the administrative work around related initiatives, such as the OSRE and graduate teaching fellowships aimed at increasing productive student engagement in open source communities.

With the Sloan Foundation grant, the new OSPO will be able to widen its open source portfolio within UC Santa Cruz and reach out to the broader UC community. The University of California—which has ten campuses and over 280,000 students—boasts many highly successful open source projects. However, it has no centralized or established method of assigning them value or effectively sharing them within open source communities. The soon-to-be-launched UC Santa Cruz OSPO aims to change that. The Center will serve as a blueprint for other public universities seeking to capture the value of successful open source projects, attract industry funding, recruit talented faculty and researchers, and better prepare students for successful careers. 



To learn more about Open Source Research Experience program, visit the [program website](#).

Papers by partners in Europe achieve noteworthy recognition

by Matej Hrušovský

Support from Red Hat Research for PhD students in France and Austria has borne its first fruits. Supported PhD students in both countries have published significant papers that have been accepted at renowned conferences: [REBLs 2021](#) (focused on reactive and event-based languages and systems) and the [Network and Distributed System Security \(NDSS\) Symposium 2022](#).

“Analysing the performance and costs of reactive programming libraries in Java,” from Red Hat engineers Julien Ponge and Clément Escoffier, PhD student Arthur Navarro, and Professor Frédéric Le Mouël (University of Lyon; INSA Lyon), was presented at REBLs 2021 on October 18, 2021. The paper discusses the performance of the three major Reactive Streams compliant libraries used in Java applications: RxJava, Project Reactor, and SmallRye Mutiny. The authors show that advanced optimization techniques such as operator fusion do not yield better performance on realistic I/O-bound workloads, and they significantly increase development and maintenance costs.

On November 16, 2021, Red Hat partners at Graz University of Technology released “[Remote memory-deduplication attacks](#).” The paper, authored by Martin Schwarzl, Erik Kraft, Moritz Lipp, and Daniel Gruss, will be presented at the NDSS conference, to be held as a hybrid event from February 26 to March 3, 2022. The work reevaluates the security risk posed by memory deduplication

Long-term support for PhD students is an essential part of the Red Hat Research program.

in a single security domain. By observing timing differences among requests, the authors executed the first fully remote memory-deduplication attacks in both Windows and Linux operating systems. Their research demonstrates that such attacks are viable risks and proposes mitigation. More information about the conference can be found on the [NDSS symposium website](#).

Long-term support for PhD students is an essential part of the Red Hat Research program. Though participants are evaluated with considerable care, it is not possible to guarantee a project’s success or its timeline for achieving major milestones. When these projects achieve recognition for their results, it is a reason to celebrate. 



About the Author Matej Hrušovský

has been with Red Hat for more than eight years, six of which have been spent managing the university program in EMEA. Aside from attracting new talent mainly from universities and schools, the core of Matej’s job is to find and put the right people from Red Hat and academia in the same room together.

Call for Code: tech for good

Solutions addressing climate change, food and water scarcity, and racial justice are among the Call for Code challenge winners.

For the last four years, IBM has issued a challenge to the world's software developers to design and build applications that would have an impact on a global environmental or social issue using open source software.

The [Call for Code](#) challenge is driven by the belief that gathering as many good ideas as possible and sharing them across the community is the most effective way to arrive at the best solution. This belief is one of the core tenets of the open source approach and its principal difference from proprietary or closed source development.

What's most compelling about IBM's Call for Code is its emphasis on social rather than technical solutions. Contestants not only must build an application in a limited timeframe, using a combination of IBM and open source technology, but also demonstrate that it solves a real-world problem. In addition, the application must be available in open source code repositories so others seeking similar solutions can use it. Access the code for all applications developed for the Call for Code challenge on Github (github.com/Call-for-Code/Project-Catalog).

Submissions are evaluated across four criteria:

- **Completeness and transferability:**
 - How fully has the idea been implemented?
 - Can it achieve an impact in the field?
 - Can it be transferred elsewhere?

- **Effectiveness and efficiency:**
 - Does the solution address a high-priority area? Does it achieve its goal effectively and efficiently? Can it scale?
- **Design and usability:**
 - How good is the design, user experience, and ease of use of the solution?
 - How quickly can it be used?
- **Creativity and innovation:**
 - How unique was the approach to solving a long-standing or previously intractable problem?

The 2021 winner is a case in point: Saaf Water, developed by a team in India to help the 2 billion people in the world without access to safe water. The project responded to the theme of the 2021 Call for Code Global Challenge, climate change. The application combines IoT and AI technology to monitor and analyze the quality of groundwater, notifying water management systems and consumers of any issues in their area. Learn more on the [IBM developer blog](#).

THE UNIVERSITY CHALLENGE

RHRQ wants to make readers aware of the Call for Code University Challenge, which accepts submissions from university students all over the world in a variety of disciplines. This opportunity is open to all students, and we hope to be reading about one of your submissions in future years.

Trashtag | 2021 University Challenge Winner *Institut supérieur d'électronique de Paris, France*

The issue of waste has severe implications for climate change, with almost 10 million tons of waste dumped into oceans and natural environments each year. Trashtag offers a platform to encourage local community members to help remove waste. Considered “Trashtag Warriors,” these community members share the work they have done through the application, using AI to analyze before and after pictures of an area that has been cleared. Jobcoin, a purpose-built cryptocurrency, is awarded through the platform as an incentive.

Agrofate *University Center FIAP, Brazil*

Agriculture is a leading contributor to climate change, especially as rainfall decreases and the water needs for farming escalate. Agrofate is tackling the problem with a simple mobile app to help small farmers get crop-watering recommendations based on real-time weather data and soil conditions. The team has seen the need for this firsthand in their native Brazil, where 35 million people don't have access to safe drinking water, and agriculture is the largest consumer of water.

FarmAid *University of Sydney, Australia*

Poor infrastructure and financial limitations cause even the smallest

flooding and drought events to affect smallholder farmers, who account for over 80 percent of the world's food supply. FarmAid provides AI solutions and advice to help farmers stay one step ahead of extreme weather events that threaten their livelihood, including access to weather reports, photo-based crop analysis, production-cycle recommendations, and more.

Mile-12 *Augustana University, NYU, PSG College of Arts and Science, US, UAE, India*

Many small farmers and food producers struggle to create a sustainable business model to support themselves and their families. This becomes increasingly difficult as climate changes and more risk is introduced into the system. This application uses AI to help food producers form cooperatives, forge links with distributors and consumers, and receive advice on enhancing profits by providing tools like price recommendations for produce.

Well Well *UTCC/Harbour Space, Thailand*

As climate change heightens the risk of drought, finding the right location for a water well is a growing problem in many parts of the world. Using machine learning and historical data, Well Well simplifies this grueling and expensive process by helping predict locations to drill. Well Well also offers a cheap portable device

that guides users to the nearest ideal place to drill based on GIS data for those without access to the internet or mobile devices.

RED HAT'S CALL FOR CODE CHALLENGE

Companies and organizations can mount their own challenges to motivate and reward their employees. Red Hat's internal challenge yielded multiple interesting applications. Some of these applications are so compelling that Red Hat associates continue working on them, and small communities are growing around them.

The winner of Red Hat's internal challenge was [Carbon Ninja](#), an application that aims to tackle carbon emissions contributed by the food industry by informing users of their carbon footprint based on a shared photo of a plate of food. The average US citizen has a carbon footprint of around 36,000 pounds a year—one of the highest in the world. While the fight against climate change requires cooperation among nations, governments, and large corporations, individual food choices are also a significant contributor. The Carbon Ninja team comprised Red Hat's Dalia Khater, Dhruv Aggarwal, Kinsey Ness, Zachary LeBlanc, and Joseph Torcasso.

CODING FOR RACIAL JUSTICE

Racial justice, especially the associated topic of technology and data bias, has generated its own Call for Code challenge.



Trashtag participant

The 2021 Racial Justice Call for Code projects winner was [Fair Change](#), a mobile application for witnesses to capture incidents and provide information to police officers aimed at avoiding escalation. The app includes a map view to enable search and visualization of incident clusters.

The complete list of finalists shows the breadth of solutions submitted:

- **TakeTwo**
An app that allows content platforms to detect and eliminate racial bias from their content
- **Truth Loop**
A platform for sharing and communicating about policies, regulations, and legislation (PR&L) in a specific community
- **Open Sentencing**
An app enabling defense attorneys to upload information about a case and client and receive analysis regarding potential bias issues
- **Incident Accuracy Reporting System (IARS)**
A content management application for capturing statements, videos, and audio feeds from first-hand individuals relating to incident reports
- **Five Fifths Voter**
A platform providing personalized information for voters facing obstacles including inexperience, mobility, and confusion about registration, early voting, and eligibility
- **Legit-Info**
An app helping citizens access policy proposals, candidate voting records, and legislation in a central database 

ABOUT THE PROJECT

The [Call for Code website](#) has complete information on how to design a winning project, how to submit it, and, best of all, how to get started contributing to an open source project.

Subscribe to Call for Code emails to learn about the 2022 theme and submission timeline as soon as they are announced. Last year's challenge launch was in late March with submissions accepted from late April to late July. Good luck with your submission!

What is Red Hat developing

NEXT?



Learn more at
next.redhat.com



About the Author
Shaun Strohmmer

is the editor of Red Hat Research Quarterly. She has worked as a writer and editor in academic publishing for over twenty years, and since 2014 she has focused on software development, cybersecurity, and computer science.



Sixteen Red Hat Collaboratory Research Incubation Award winners announced

Funding recipients will study AI in cloud operations, hardware stack innovations, performance improvements, and more.

by *Shaun Strohmmer*

The Red Hat Collaboratory at Boston University recently announced the recipients of its first-ever [Research Incubation Awards](#). Reviewers from BU faculty and Red Hat selected sixteen proposals to fund, including one large-scale, multi-year project and another five smaller-scale projects that pair BU faculty and graduate students with Red Hat engineers. If you are interested in exploring collaboration opportunities on these or other research projects, please contact [Heidi Picher Dempsey](#), US Research Director for Red Hat.

Orran Krieger, BU professor of electrical and computer engineering and co-director of the Red Hat Collaboratory, said of the award winners, "I am excited not only about many of the individual projects, but how they can build on each other. Fundamentally, we want to get everyone out of their comfort zone to do things that are only possible because of the collaboration. Researchers tend to work

on projects where they can control all the variables, and it is difficult for others to build on the research artifacts developed by graduate students. Engineers are often focused on the next incremental change and rarely have time to use state-of-the-art research. This collaboration not only enables research projects that have an impact but also enables research artifacts that are of high enough quality that other research projects can build on them."

Faculty members Ayse Coskun, Alan Liu, and Gianluca Stringhini; Red Hat researchers Steven Huels, Marcel Hild, and Daniel Riek; and IBM researcher Fabio Oliviera won the award for a large-scale project and will receive \$1 million in funding over two years. Their project, "AI for Cloud Ops," aims to deliver easily accessible, open source AI technologies for developers and administrators who are solving real-world performance, resilience, and security challenges. The researchers will create new methods for fusing and representing systems



Pictured from left: *Gianluca Stringhini, Alan Liu, and Ayse Coskun*

data to enable AI-based analytics and will build, apply, and scale AI frameworks to improve performance, management, security, compliance, and resilience problems in the cloud.

These are the winners for small-scale proposals:

- Faculty members Christos Cassandras, Vasiliki Kalavri, John Liagouris, and Mayank Varia and Red Hat researchers Alexandra Machado, Jim Craig, and Christopher Tate will investigate the link between wellbeing and eco-smart cities in the project "Creating a global open research platform to better understand social sustainability using data from a real-life smart village."
- Faculty member Martin Herbordt and Red Hat researchers Uli Drepper and Ahmed Sanaullah will create a generic operating system and firmware to standardize software interfaces and increase code compatibility in the project "DISL:

A Dynamic Infrastructure Services Layer for reconfigurable hardware."

- Faculty members Vasiliki Kalavri and Jonathan Appavoo and Red Hat researcher Sanjay Arora aim to prove that energy efficiency and the layers of software in open source stream platforms are not incompatible in the project "Towards high performance and energy efficiency in open source stream processing."
- Faculty members Manos Athanassoulis and Renato Mancuso and Red Hat researchers Uli Drepper and Ahmed Sanaullah are developing a hardware-software design methodology for data systems that implements near-memory processing in the project "Near-data data transformation."
- Faculty member Martin Herbordt and Red Hat researchers Uli Drepper and Ahmed Sanaullah are creating an open source tool for reducing application development effort

and turnaround time for Field Programmable Gate Arrays (FPGAs) in the project "Practical programming of FPGAs with open source tools."

Ten awards also went to faculty projects related to speculative, fundamental research or Mass Open Cloud projects designed to initiate a collaboration:

- "Privacy-preserving cloud computing using homomorphic encryption," Ajay Joshi
- "Symbiotes: A new step in Linux's evolution," Jonathan Appavoo
- "Foundations in open source education," Jonathan Appavoo
- "Secure cross-site analytics on OpenShift logs," John Liagouris
- "Robust data systems tuning," Manos Athanassoulis and Evimaria Terzi

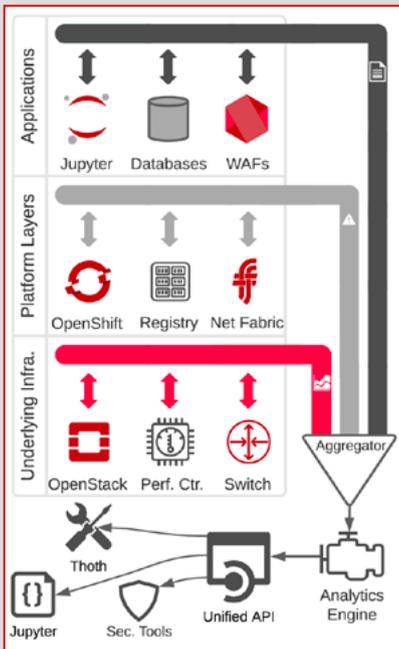
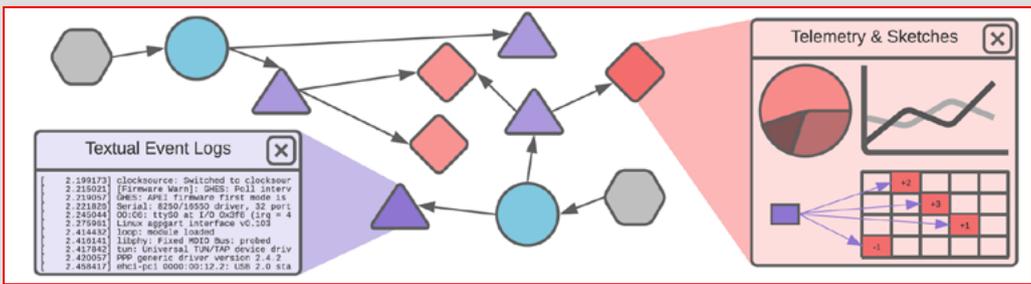


Figure 1: An illustrative overview of the AI for Cloud Ops project, which aims to demonstrate the performance, resilience, and security benefits of AI-driven cloud analytics in modern continuous integration/continuous deployment environments. The project will make customized analytics available to developers and administrators via queryable APIs during open source software deployment (e.g., through Jupyter notebooks) and at runtime.

Figure 2: The cross-layer data visualization that will be designed as part of the project. Each geometric shape in the trace graph represents a cloud entity, such as a container, service, load balancer, etc. A key distinguishing feature of the project is that it will design cross-layer analytics frameworks, starting from the application layer and connecting to deeper system-level instrumentation.



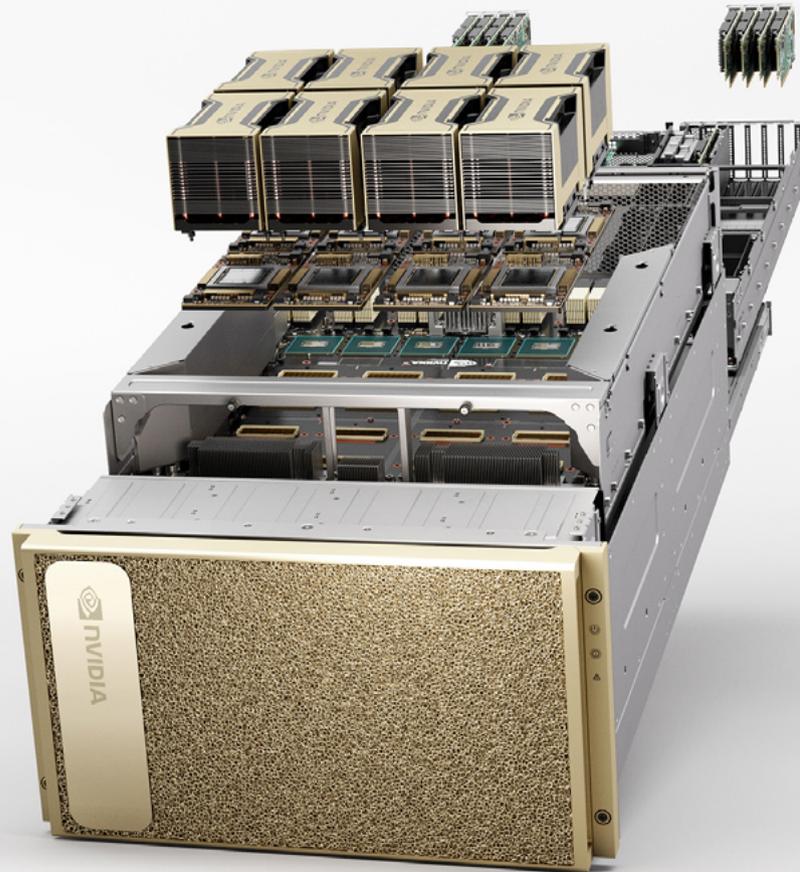
Awards continued...

- “OSMOSIS: Open source multiorganizational collaborative training for societal-scale AI systems,” Eshed Ohn-Bar
- “Serverless streaming graph analytics,” Vasiliki Kalavri

- “Enabling intelligent in-network computing for cloud systems,” Zaoxing Liu
- “Linux computational caching,” Jonathan Appavoo
- “Intelligent data synchronization for hybrid clouds,” David Starobinski



Learn more by visiting research.redhat.com to find individual project pages. Please stay tuned to Red Hat Research for other opportunities to learn about these projects and get involved.



THE UNIVERSAL AI SYSTEM FOR HIGHER EDUCATION AND RESEARCH

NVIDIA DGX A100

Higher education and research institutions are the pioneers of innovation, entrusted to train future academics, faculty, and researchers on emerging technologies like AI, data analytics, scientific simulation, and visualization. These technologies require powerful compute infrastructure, enabling the fastest time to scientific exploration and insights. NVIDIA® DGX™ A100 unifies all workloads with top performance, simplifies infrastructure deployment, delivers cost savings, and equips the next generation with a powerful, state-of-the-art GPU infrastructure.

Learn More About **DGX** @ nvidia.com/dgx-pod

Learn More About **DGX on OpenShift** @ nvidia.com/dgx-openshift

News



Open source researchers in security and education win 2021 innovation awards

by Vashek Matyáš

About the Author

Vashek Matyáš is a professor with the Centre for Research on Cryptography and Security at the Faculty of Informatics, Masaryk University.

Masaryk University (MU) awarded the 2021 MUNI Innovation Award at its Business Research Forum on November 11, 2021. These awards recognize individual students and staff whose research has been successfully implemented in practice, helped to improve products or services, or in some other way enhanced the social relevance of MU research. Awards went to seventeen individuals or teams from ten faculties, including three to Faculty of Informatics researchers who have worked with Red Hat Research on advanced open source projects.

Petr Švenda received an Innovation Award for his design of new methodologies for security system analyses and reviews of cryptographic implementations. These methodologies enable verification of the security of newly developed devices and find bugs, allowing for their timely correction against misuse. In 2017, his group found and helped eliminate the

most serious cryptographic vulnerabilities in many European countries' current electronic ID cards and in the security solutions of major IT vendors. The results of this research led to a change in global certification procedures in the field of digital signature key generation. Petr presented his recent work on certification at a Red Hat Research Day in March 2021 ("[Mining issued common criteria and FIPS 140-2 certificates – more transparency for developers, vulnerability researchers, and society](#)").

Radek Pelánek won the MUNI Innovation Award for research results using machine learning and data-processing methods that provided the basis for the design of the learning environment implemented on the website [umimeto.org](#). Radek leads the university's Adaptive Learning Research Group, which focuses on using artificial intelligence, machine learning, and data analysis in the development of learning environments.

Umíme is a startup founded by Petr Jarušek, a former PhD student in the group. Umíme started with orthography and grammar and now covers many other subjects, including mathematics, English, programming, and geography. The site is used by more than ten percent of Czech schools. Red Hat currently supports his research with Tomáš Effenberger, a PhD student focusing on introductory programming and computational thinking. Their research concerns the design of exercises (e.g., [microworlds for block-based programming](#)) and analysis of data from programming exercises.

Milan Brož, Mikuláš Patočka, and Vashek Matyáš were awarded for their work on Full Disk Encryption (FDE), which has become a widely used security feature. Until now, FDE rarely provided cryptographic data integrity protection; the team introduced an algorithm-agnostic solution that provides both data integrity and confidentiality protection at the disk-sector layer. Their open source solution is intended for drives without any special hardware extensions and is based on per-sector metadata fields implemented in software. Their implementation has been included in the Linux kernel since version 4.12. Milan did this work as a Red Hatter in the Brno office and as a PhD student of Vashek Matyáš, with Red Hatter Mikuláš Patočka contributing mainly to efficient implementation of the solution.

[Work on FDE](#) continues through the cooperation of Red Hat Czech and MU. 

RH
RQ

Read current articles

Browse back issues

Subscribe to print or digital

research.redhat.com/quarterly



Red Hat 
Research
Quarterly

Clouds

on the horizon

*Shared cloud computing
resources make research
more accessible
and more powerful*

Interview by **Heidi Picher Dempsey**

Interview

Dr. Michael Zink is Professor of Electrical and Computer Engineering at the University of Massachusetts, Amherst, with research interests in sensor networks, next-generation internet, systems engineering, and multimedia distribution. In addition to publishing and teaching, Dr. Zink has participated in several projects providing distributed systems and virtual networks for research and education, including GENI and ExoGENI (2007-2021), Cloud Lab (2014-2021), and now the Open Cloud Testbed (OCT) since 2019. The OCT, a collaboration among researchers from Boston University, Northeastern University, and UMass Amherst, was recently awarded \$5 million from the National Science Foundation (NSF) to develop a testbed for new cloud computing platforms, combining research and production cloud capabilities in shared testbeds, as well as new features such as programmable FPGAs for cloud developers.

Heidi Picher Dempsey: I want to start with your interest in sensors and your secret history as an electrical engineer, which some may not know. How did you start there, and how did that grow into cloud-related projects?

Michael Zink: After graduating from high school, not knowing what to do, I thought electrical engineering seemed like fun, so I started as an undergrad in electrical engineering. Actually, all my degrees are in electrical engineering. But electrical engineering and computer science are not so far apart. For example, when I was an undergrad in Darmstadt, Germany, we had a lab where we operated a radar—I think that was my first foray into sensors. When I graduated with my PhD, I looked for a postdoc position. I got one here at UMass, in the NSF Engineering Research Center for Collaborative Adaptive Sensing of the Atmosphere (CASA) that does low atmospheric weather observations.

We built a closed-loop weather observation system. At that time, I had worked on multimedia streaming for my PhD, so I got a lot more expertise in distributed systems and networking, but I also understood the engineering side. I was considered someone who fits in nicely because I knew about atmospheric sensors, and I knew how data was shipped around and processed.

That was a beautiful experience for me because I learned to work with many principal scientists.

That's what I still do today: make those connections between scientists' work and their compute needs. I just came back from lunch with two professors from the physics department speaking about the cloud and using [ESI \(Elastic Secure Infrastructure\)](#). The BU folks and the UMass folks are providing an upgrade for the ATLAS NET2 node, which does the data processing and distribution for the Large Hadron Collider in Switzerland. The idea is that we will use ESI to make this hardware available not only for their compute needs, but for other researchers also.

Heidi Picher Dempsey: Let's back up and [explain ESI](#) for those unfamiliar with it. ESI is one of the collaborative projects on the OCT. It aims to make it possible for datacenters and research groups to share bare metal machines. As demand fluctuates in different research centers, you can share resources. You won't find researchers running out of compute time, memory, or network to do their work, which unfortunately happens now.

Michael Zink: Exactly.

Heidi Picher Dempsey: OK—back to your work in radar. You looked at the amount of land you could



About the Author
Heidi Picher Dempsey, US Research Director for Red Hat, and Dr. Zink are long-time colleagues who originally met working on the NSF GENI project. She interviewed Dr. Zink about his research and teaching career, the promise of making computing more accessible and efficient, and the social impact of science.

cover with one of these radars, and you had to guess where to put them or move them based on the data you were getting. Are you still looking at those kinds of problems? How involved do you stay in that research while you're starting all this other work?

Michael Zink: The new weather-observation area I've been looking into is drones. Unpiloted aerial vehicles are becoming more and more popular. What we hear from people working in this environment is, "When there's weather, we don't fly." I'm not an economist, but I think that's not a good business model. We've been spending a lot of time in path planning based on weather observations with our systems. We're working with folks from the business school to discover if, when you have certain stochastic information, you can carry out a flight successfully.

Information has to become available much more quickly because drones often fly shorter distances. You don't have the luxury of flying large diversions, as you can with a Boeing 747. We've been looking into that issue, until recently, using a lot of the ExoGENI infrastructure. Now we are looking into the [Chameleon testbed](#), an OpenStack-based community testbed hosted by the University of Chicago and the Texas Advanced Computing Center at UT Austin. Especially with Chameleon at the edge, we hope to see how edge-to-cloud architectures, with the compute resources provided on that spectrum, can make these applications possible.



Mike Zink with students

Heidi Picher Dempsey: Edge computing has gotten a lot of attention recently, but you've been working on it for a very long time.

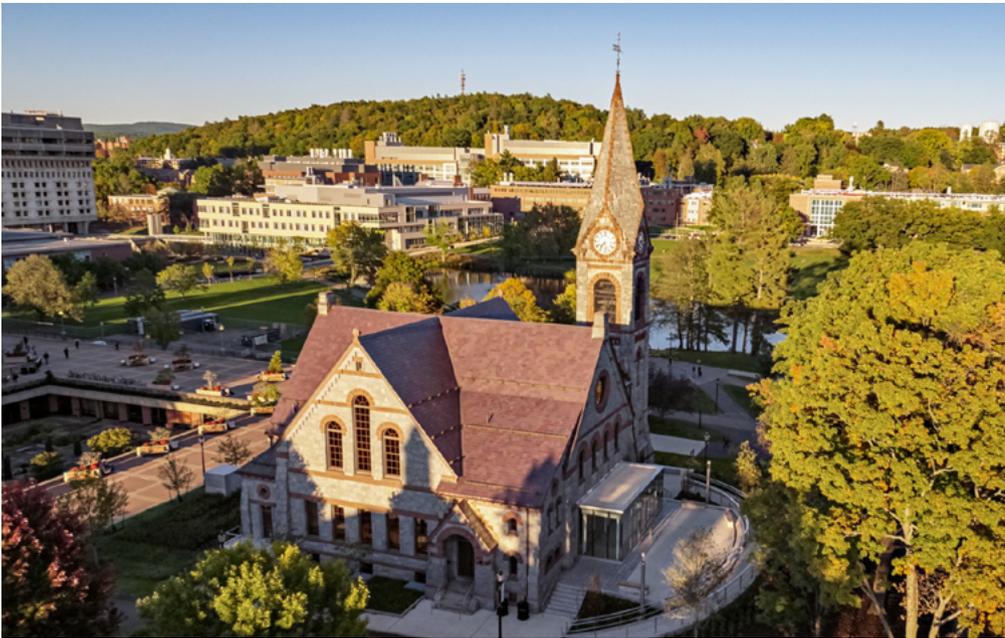
Michael Zink: We were forced to do this with radars because we had high data rates. We were compressing this data then sending it from a remote radar to a location where we could do more processing. Now, with the Internet of Things (IoT) and connected autonomous devices, we see that some processing has to happen at the edge, because there's no way you can ship all this data back to a central location. That will be even more important in the future.

Heidi Picher Dempsey: You mentioned ExoGENI and the NSF GENI program,

which started more than ten years ago. So you had the experience of being involved in a project at the beginning and building an infrastructure that didn't yet exist, to provide high data processing rates, very flexible networking, high throughput, and dynamic composition. All of these things are now in the commercial cloud. With the recent shutdown of ExoGENI, you've seen the end of an infrastructure project, but all of the research topics have not been solved. How does it feel to spend that many years on a research effort but then not be able to finish?

Michael Zink: That's life, I guess. <laughs>

Heidi Picher Dempsey: Philosophy in computer science!



University of Massachusetts Amherst

Michael Zink: In most sciences, right? You always have very high goals, which is good because that motivates people. If we achieve all these goals, that's more than perfect—we will not always be able to do that. But if many good discoveries come out along the path, that's wonderful.

For me, it has been an incredible experience, being there from the beginning, seeing these crazy ideas. Some didn't materialize, but others, which you wouldn't think were possible, did. And the human network we built around it may be the biggest outcome. The research is important, but so are the students who had the experience. If I look at my students who are now at Apple or Akamai, bringing their expertise back in these companies—

that are not necessarily cloud companies, but they're using all the same approaches—that's an incredible impact. That's just two examples from maybe thousands of students who went through that project at the time. To a certain extent, that's why we see some of these technologies now as commercial offerings.

Heidi Picher Dempsey: And if you're coming up with questions that generate more questions, you're almost guaranteed never to get to the end, if it's a nice rich project with a lot of meat to it. A lot of the value is in the questions, which is hard to get across to new students.

Michael Zink: Oh, yes. Sometimes it's hard to keep students focused on

their initial question because, once they get going, they have so many more questions. Sometimes you have to say, "You have to answer at least one question before you go to the next one."

Heidi Picher Dempsey: That brings us to another thing I wanted to ask you about. I've always admired your teaching style and how you get students working on real systems. As people are experimenting with infrastructure that gets more and more complex, have you found ways of dealing with that complexity and getting people started more easily? Or is it still one-on-one mentoring that makes the difference?

Michael Zink: Thank you—I wish all my students would've listened to that first part! That's a struggle. There's so much knowledge about certain software, systems, and technologies that we have accumulated over time. Even if they are the smartest people in the world, students can have a hard time catching up. The approach I take is to limit the set of tools they need. It's not totally under my control, but if possible, I'd have them use a single testbed throughout their career as a graduate student. If they use five testbeds, they have to learn all the details about all five and spend all their time on that.

But what do you do when things change, like ExoGENI going away? The funding agencies have to help by ensuring that the new testbeds are not radically new. This is a little bit of an oxymoron, right? We always want to have the latest and greatest.



Mike Zink lecturing at Red Hat Research Day Europe 2020

But from a grad student perspective, some continuity is good. We all have to be a little bit creative. For example, [FABRIC](#) is a new large-scale research infrastructure that some see as a follow-up to GENI. Could there have been virtualized GENI slices in FABRIC that work like the GENI technology we had in the past? If we had done that, students would have had an easier life, and they could have transitioned over time.

Heidi Picher Dempsey: Something that might help, too, is if the tool they interface with stays the same. You are one of the early users of Jupyter Notebooks to help students interact with the infrastructure. Do you think that kind of tool is helpful and likely to stay around?

Michael Zink: I've often been proven wrong, so I hate to make predictions, but these tools are amazing. At UMass, we run a cluster that has Jupyter Notebook as a front end. You can put an undergrad in front of that cluster and have them use multiple GPUs, whatever they want, to do computing they could never do in the past.

This is a unifying thing. I teach a sophomore class right now and if they need more resources than they have on their laptops, they can use Jupyter Notebook and do stuff there. They can go on our cluster. That is a standard that becomes more and more important. We see people from the social sciences now using that cluster and using Jupyter Notebook because it's more accessible.

People who you'd never think of using a cluster ten years ago are using it because it comes with an interface that makes it much easier to use.

Heidi Picher Dempsey: That's a really exciting trend. We're starting to see that in Boston University projects, too: We can reach out beyond computer science, math, and physics to people in other fields who can use the compute power. Do you think this will make students less interested in Linux in the long run because it's now further down the stack?

Michael Zink: I hope not, because we still need those students!

I hope it will just cause more curiosity. I'm sure there are already high schools using Jupyter Notebooks. After getting them in touch with this technology, maybe there's five or ten percent who want to know what the underlying mechanisms are and get into Unix and Linux, even down to architecture in some cases.

Heidi Picher Dempsey: You can send them all our way!

Michael Zink: No, no, no! I need them myself.

Heidi Picher Dempsey: Let's talk about your recent award from NSF for the OCT project. Why do you think that project's important, what do you think attracted NSF's interest, and what do you hope to do?

Michael Zink: We have to make clear that this is for cloud computing research,

right? A researcher with a need for compute resources can often use those on campus, or they can get it through the three big providers—that’s perfect. But those are closed environments. We don’t know the inner workings of Amazon cloud or Google cloud, for example. We want to provide testbeds that allow researchers to run their own software, all the way down to the operating system, and have much more freedom on the networking side. In some instances, we also have technology that’s not been made openly available. For example, the field-programmable gate arrays (FPGAs) we put in the OCT are something we see a need for in the research community.

The sharing aspect is vital. We cannot afford to give every researcher their own testbed. My goal is to make a testbed available where researchers can go down to the hardware and, without damaging anything, do what they need to do to perform the research, have control, and make it reproducible. This is a kind of stigma we have in our sciences compared to the natural sciences: it’s hard to repeat an experiment. OCT is providing these mechanisms and helping the community to do the work they want to do.

Heidi Picher Dempsey: When you’re trying to give people as much flexibility as possible, how do you also consider security?

Michael Zink: For a research testbed, you have to have a certain level of freedom, but you don’t want to end

up on the front page of the Boston Globe because someone used it to mine Bitcoins. One way to achieve that is through authentication and identity management. The community, with the help of Internet2 and industry, has done a great job working on this.

I think we have reasonable mechanisms in place. Thanks to the work Orran Krieger did at BU, this is on our minds. If we share, for example, bare metal servers, how can we make sure a new researcher gets an uncompromised system? We’ve been forced to think about this already with the FPGAs. If you reboot the host, you don’t necessarily reboot the FPGA, so you might leave stuff behind for the next user. We’re working hard on ensuring that every time a new user comes on, we’ve scrubbed the FPGAs, so there’s nothing that either shouldn’t be exposed to someone else or could cause any problems. That’s important because if people don’t feel a certain level of security is provided, they won’t use it. They need to feel confident.

Another struggle is preserving the data. You’re working hard on an experiment, collecting all this data, then suddenly it’s gone. You have to provide mechanisms or at least awareness about where the data resides, what you have to do to secure data, because that upsets people a lot.

Heidi Picher Dempsey: Exactly—and upset is a nice word some days! You and I have both lost data. It’s a delicate balancing act to allow people to change the things they need to change, to

Datacenters use an incredible amount of resources. This has a societal impact, for example contributing to climate change. We have to continue to make computing more power efficient.



explore things and ask new questions and answer them, but still make sure that all the people sharing the testbed don't stomp on each other.

Michael Zink: It comes back to education. When students use the testbed for the first time, we have to tell them that their actions can have consequences. They need to think a little bit before pushing a button and shutting down the internet.

Heidi Picher Dempsey: One of the things we are trying to do in the [Mass Open Cloud](#) is to make it possible for people to collect and share data about how the systems are working—basically, open telemetry. Is that one of your goals for the OCT too?

Michael Zink: One thing we did already for the Cloud Lab project is measuring the power consumption of systems. Datacenters use an incredible amount of resources. This has a societal impact, for example contributing to climate change. We have to continue to make computing more power efficient. By making data available, when a researcher runs an experiment, they can say, "If I run it this way, I know I use so many kilo or megawatt hours. If I run it that way, I use thirty percent less."

We want to make this available in the OCT also, and it's relatively straightforward to implement. We don't have much influence over how computation is happening, but we can observe and make data available as much as possible.

Heidi Picher Dempsey: How do you deal with recording the hundreds of variables that might be affecting that data at the time that it ran? For example, the versions of all of the pieces of the stack that were running at the time

Michael Zink: It's a nightmare!

Heidi Picher Dempsey: Are you going to solve a nightmare for us?

Michael Zink: This has been a topic forever. It's easy to say, "This server uses so much power at this instance." OK, but what was it running? What operating system? What CPU? How do you collect all this information? That's a tough problem that—I'm sorry, don't tell anyone, but I'm not going to solve it.

Heidi Picher Dempsey: Well, we have to leave a few challenges for everyone else.

Michael Zink: Yes. Let's let someone else do that!

Heidi Picher Dempsey: I agree, though, especially when you include the application levels of the stack, where the data is private, it's incredibly challenging.

Michael Zink: I would add one small thing: It's important to think about the minimum set of information I can live with to do my analysis. We often think the other way: I'll take as much data as I can. That's not good. Other sciences work so hard

on measuring just one variable. We should be smart and consider when less is more. Getting the data is always easy; maintaining it and making sense of it is the biggest problem.

Heidi Picher Dempsey: I have a fun fact about you: you're one of the few computer scientists I know who's been in a commercial.

Michael Zink: What?!

Heidi Picher Dempsey: There were commercials advertising UMass, and they included a view of what your students were doing and a few words from you. You don't remember this at all?

Michael Zink: I don't remember this! I'm so sorry.

Heidi Picher Dempsey: Well, we can see that Professor Zink is completely dedicated to his academic goals and cares not for the rest of the world.

But you do have to promote your work to attract students. I used to work at the Woods Hole Oceanographic Institute, and Bob Ballard, a famous scientist at Woods Hole, had a unique point of view on this. He said, "I love science. I've dedicated my life to science. But I don't get a kid sitting in their bedroom excited about science by discussing it on an abstract level." So he made a great effort to make the underwater remotely operated vehicles something students could interact with. Think of how huge of a challenge that was in the 1990s. Making sure to capture the imagination of others was just as important as the science he was doing.

Michael Zink: For a big child like me, that's super exciting. In the Massachusetts Green High Performance Computing Center

(MGHPCC), our mission is to provide compute resources for scientists coming up with new discoveries. For example, today I met with the guys in physics who crash particles into each other to make discoveries all the time. There are structural biologists who have crazy microscopes to see molecule structures.

You can find shipwrecks at the bottom of the ocean at depths that used to be unimaginable, find new forms of life. That's the essence of it. The Resource Reservation Protocol (RSVP) work I did many moons ago was researching allocation to stream video in a better way. Now we have Netflix and Disney+ and all this streaming content. It's for the greater good, I hope, because it entertains a lot of people.

Heidi Picher Dempsey: It's like we built the matchstick. You need that other person to light it.

Michael Zink: That leads to another important point. I'm talking to two women for this interview today [RHRQ editor Shaun Strohmer also attended the interview], and that's awesome. But the demography of our field is not that diverse. Studies have shown, and we hear on campus a lot, especially from minorities, that students want to contribute something for their peers.

If they become engineers, they want to be able to do something that benefits their community as well. We often miss that in our message. Yes, it's important to figure out how black holes are composed, but that can be too abstract. It's more important that people in a food desert can find a store with good, affordable food.

The work we do provides tools and resources that have a far-ranging societal impact. That can only grow as we can make them available to more types of researchers and more communities. 

Feature



About the Author

Daniel Bachar

is an MSc student in the Efi Arazi School of Computer Science at Reichman University. Daniel is fascinated by distributed systems, networking, data streams, and machine learning. His research focuses on the optimization of distributed container-based systems and service mesh, specifically Kubernetes.

Optimizing Kubernetes service selection

Is there a way to implement load balancing in multicloud environments that won't increase resource usage?

New research suggests the answer is yes.

by *Daniel Bachar*

Multicloud providers and microservice-based applications across clouds are becoming increasingly popular. Organizations that use them enjoy the benefits of high availability, performance improvements, and cost effectiveness.

However, as microservices communicate with each other over a network, the deployment of microservices-based applications in multiple clouds runs the risk of increasing the application execution time, traffic cost, and more. For example, two dependent services tend to communicate frequently. Allocating them on different clouds leads to high cost and high delay.

While working on my thesis with Prof. Anat Bremler-Barr from Reichman University and Prof. David Hay from the Hebrew University, we were contacted by Red Hat Research team lead Idan Levi and engineer Livnat Peer. They introduced us to an interesting problem: how to optimize the service selection mechanism given a certain service distribution among multicloud and multicloud deployments. The load-balancing approaches used thus far have not been adequate for a multicloud environment for several reasons, explained below.

Our research suggests a new approach for addressing the problem. We have introduced a new system called KOSS (Kubernetes Optimized Service Selection), which optimizes service selection specifically for Kubernetes. KOSS utilizes [Submariner](#), which is used to connect Kubernetes clusters networking layer (across geolocations and cloud providers), forming what is called a multicloud or a cluster set. Submariner forms a full L3 network mesh connection in Kubernetes multicloud environments, using encrypted VPN tunnels, and allows cross-cluster service discovery.

THE SERVICE SELECTION PROBLEM

Prof. Bremler-Barr, Prof. Hay, and I defined the following service selection problem and model. When optimizing the communication load between microservices, there are two types of communication layouts to be considered: the inner-cluster (i.e., between services in the same cluster) and the inter-cluster (i.e., between microservices located in different clusters). We focus on the inter-cluster layout and offer a solution for the optimization problem derived from it. For each layout, the system must attend to two main questions.

One: How and where to implement the load balancing

Traditionally, there are three load-balancing architectures:

- **Instance-oriented**, where a central load balancer is aware of all the service instances and balances the load between them. While this approach can balance loads across instances well, it introduces price and performance overhead when communicating between the load balancer and the instances.
- **Microservice-oriented**, where services of the same type are grouped together and assigned a load balancer that directs the traffic between the instances of the same type. While this approach may alleviate the overhead introduced in the instance-oriented approach when applied within the same cluster, it does not take into account service-to-service interdependency, and it will introduce communication overhead when used in a multicluster environment.
- **Client-oriented**, where the client itself uses a load balancer (client here refers to an instance of a microservice). The load balancer is usually implemented as a sidecar application that shares namespaces and resources with the application container and acts as a proxy for the requests. This approach is widely used in the industry and adopted by all service mesh solutions. The approach offers resilience and fine-grained control for balancing the load

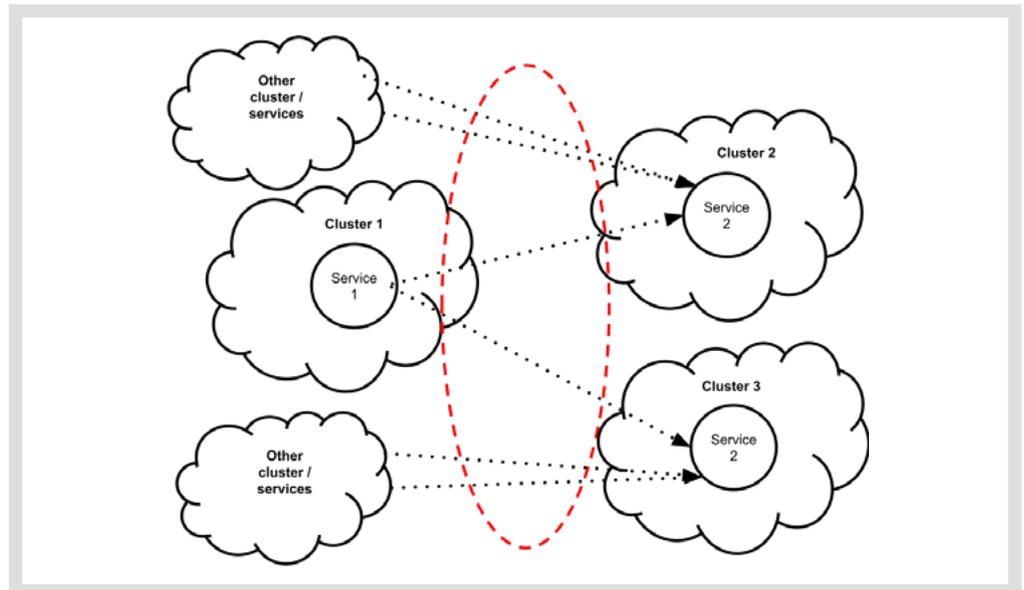


Figure 1. One-hop optimization

in the system. However, it introduces two additional hops between the instance and the sidecar application. In addition, the sidecar application requires extra resources for operation and propagation of the policies.

We suggest a new approach: **cluster-oriented**, where the selection decision is made in a centralized component and yields policies that are updated and distributed at each cluster for local and efficient routing.

Two: How to optimize the service selection

Optimization-wise, there are two main concerns:

1. **Optimization parameters:** latency, pricing, congestion, response time, etc.

2. **System constraints:** capacity, load, special topology knowledge (such as one-hop, where the model optimizes only one hop between two services, as in **Figure 1**, and chain-based, where the full dependency graph of the application is taken into account), etc.

The service selection optimization problem lies at the heart of our system. Our system and the corresponding optimization problem are working in a hop-by-hop manner, where only the next-hop destination of each request is considered for service selection. Nevertheless, we handle the interdependence of services over every single hop and accommodate competition between the different services, as shown in Figure 1.

We summarize the notation we use in **Table I**. We define the following optimization problem:

- (1) *minimize* $\sum_{c \in C} \sum_{s \in S} n_{c,s} * cost(c, s)$
- (2) *subject to* $\sum_{s \in S_t} n_{c,s} = N_{c,t}$ for all $c \in C, t \in T$
- (3) $\sum_{c \in C} n_{c,s} \leq cap(s)$ for all $s \in S$
- (4) $\sum_{c \in C} n_{c,s} \geq \alpha(c, s)$ for all $s \in S$

The optimization problem (1) aims to minimize the cost of sending requests within the system, while (2) prevents requests from being dropped, (3) avoids suppressing the instance capacity, and (4) allows minimal traffic per service instance. We defined an integer linear programming (ILP) optimization and solved it. With the solution—that is, the assignment for each variable the optimization problem will yield—we calculate a weight for each specific cluster to service instance relations as follows:

$$(5) \quad w_{s_t^c}^{c'} = \frac{n_{c',s_t^c}}{N_{s,t}}$$

Because we use weighted-round-robin load balancing, which supports floats, we allow the relaxation of the constraints and treat the variable as continuous.

KOSS SYSTEM DESIGN

While working on the implementation of the KOSS system, I had great help in designing the system and implementing it from the Submariner and multicluster development team, including Vishal Thapar, Mike Kolesnik, Tom Pantelis, Sridhar Gaddam,

Table I - Key notation summary	
C	Set of clusters (i.e., location)
T	Set of service types
S	Set of service instances
S_t	Set of service instances of type t $S_t \subseteq S$
S_t^c	A single and unique service instance in some cluster $c \in C$ of type $t \in T$
$cap(s_t^c)$	Capacity of service instance S_t^c
$cost(c', s_t^c)$	Abstract cost function describing the penalty for sending traffic from cluster $c' \in C$ to service instance S_t^c
$N_{c,t}$	The amount of requests cluster c wish to send to service type t
n_{c',s_t^c}	The amount of requests cluster c is assigned by KOSS to send to service instance S_t^c

Miguel Angel Ajo, Stephen Kitt, and others, led by Nir Yechiel. Our cluster-oriented approach, which separates the control and data planes and allows efficient updates while supporting fine-grained control over the load balancing, is described in **Figure 2**. KOSS seamlessly uses Submariner control-plane components and forms its own distributed control-

plane (the Broker) and data-plane ([Lighthouse DNS plugin](#)) layers. In a nutshell, the broker, which has a global view of the system, solves the optimization problem and distributes its outcome to other clusters. The Lighthouse plugin, a DNS service discovery, makes efficient and independent routing decisions while collecting metrics at each cluster.

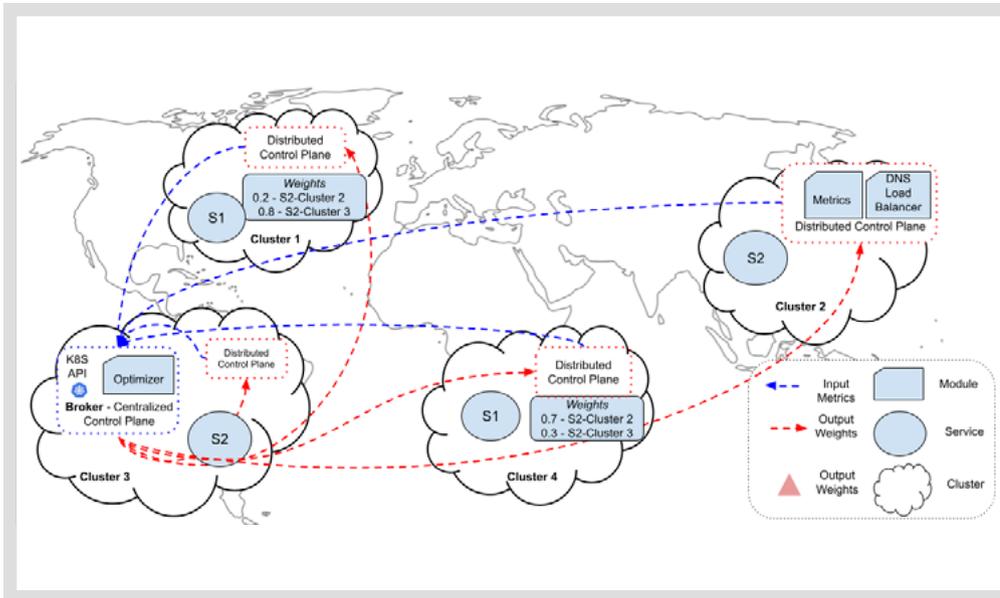


Figure 2: KOSS implementation over Submariner control plane. A centralized broker running optimization and decentralized Lighthouse plugin runs metrics collection and load balancing at each cluster.

EXPERIMENTS AND EVALUATION

In the evaluation, we take into account two metrics for the cost function:

- Price, which we define as the egress traffic price per GB of data
- Latency, which we define as the ping latency between two clusters, measured in milliseconds

We specifically chose to optimize latency and pricing because these two parameters are crucial for most operators and easy to measure. We use the **weighted sum** technique to normalize the chosen parameters for the cost function and support the multi-objective optimization.

$$(6) \text{ cost}(c_i, t_i) = \frac{p_{i,j}}{\max(P)} * \beta + \frac{l_{i,j}}{\max(L)} * (1 - \beta)$$

For example, when optimizing requests sent from cluster c_1 to service type t (which has instances deployed in cluster c_2 and c_3), we decided to allow the same importance for both the price and the latency parameters. Hence, we defined $\beta = 0.5$.

Experiments

We conducted numerous experiments and simulations, and [the simulator and dashboard](#) we have built are available online for everyone to use. We outperformed the current round-robin implementation and avoided herd behavior and underutilization. Across different experiments, we showed up to an 85% latency improvement and a cost reduction of egress traffic by up to 29%.

We suggest a new approach: cluster-oriented, where the selection decision is made in a centralized component and yields policies that are updated and distributed at each cluster for local and efficient routing.

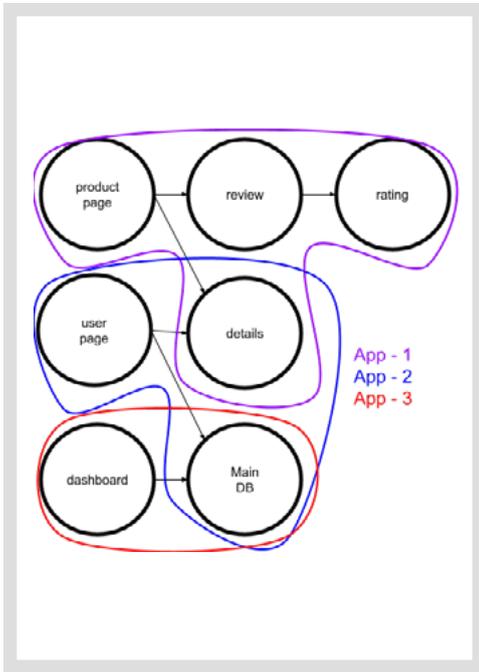


Figure 3: Multiple applications sharing services dependency graph

The following is a relatively complex case that shows our system responsiveness, flexibility, and performance.

Several web applications with microservice interdependence:

We deployed a mix of web applications that share common services, which causes competition between the applications over resources in the system. For example, we can see that both the *dashboard* service and the *user page* service are dependent on the *main DB* service. We show how our solution accounts for the interdependence between different service chains. We deployed the services on five clusters spread across five different geolocations (Brazil,

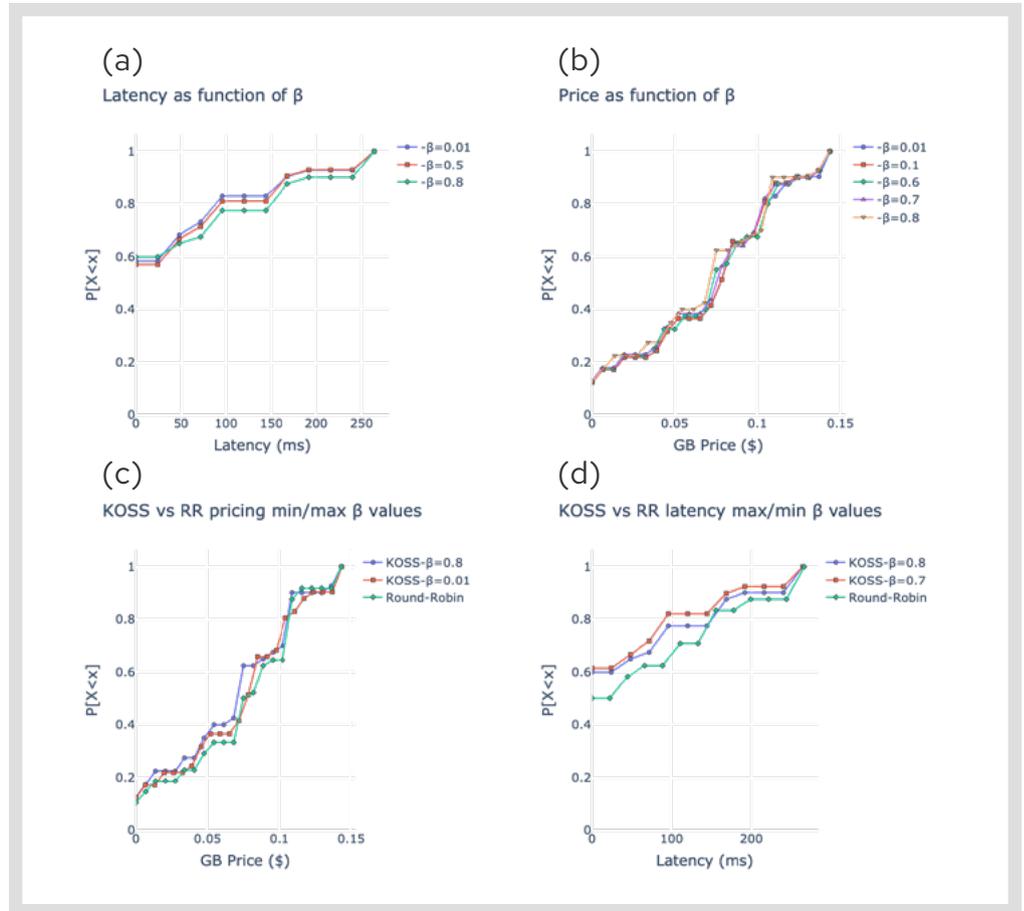


Figure 4: Multiple applications sharing the same services, (a) CDF of price as function β , (b) CDF of latency as function of β , (c) CDF of pricing as function of β , KOSS vs. round-robin, (d) CDF of latency as function of β , KOSS vs. round-robin

Japan, Germany, US East and West) and over two cloud providers, Amazon Web Services (AWS) and Google Cloud Platform (GCP). **Figure 3** is a dependency graph of the different applications. **Figure 4** is a series of cumulative distribution function (CDF) graphs: **Graphs a** and **b** show the shift in the mean metric (price/latency) according to the change in the β parameter (6). Our technique

clearly provides fine-grained control over each parameter we defined for the cost function. **Graphs c** and **d** show our superiority over the current round-robin implementation. We show that no matter what the β value we choose, KOSS surpasses the round-robin and improves the mean latency and price per GB. In this layout, we achieved a total cost reduction of up to 7% and a latency improvement of up to 28%.

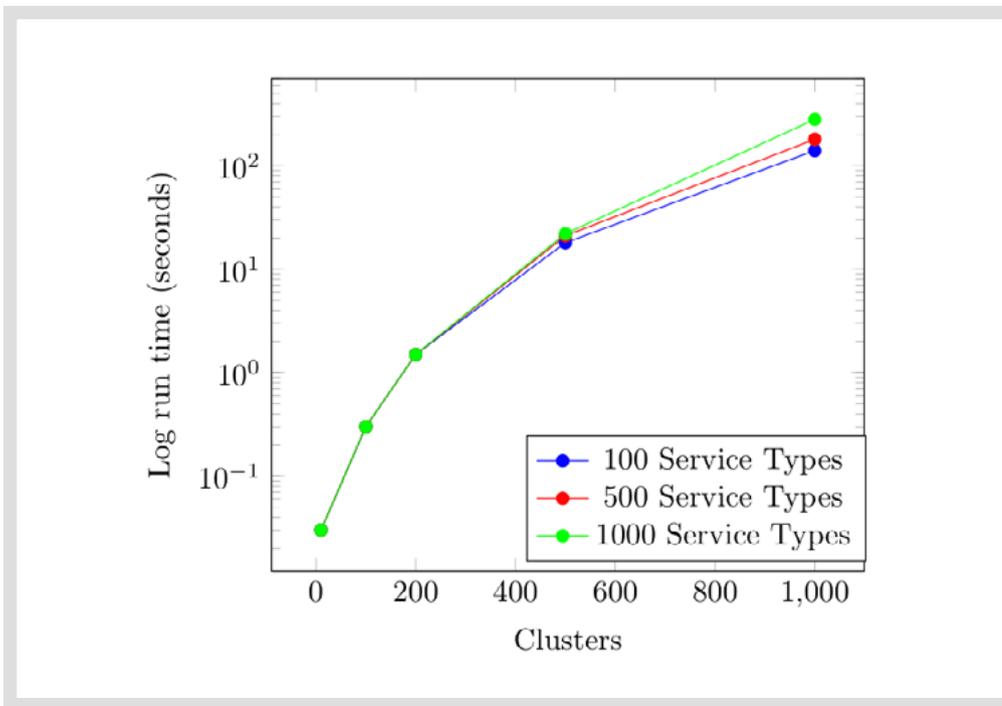


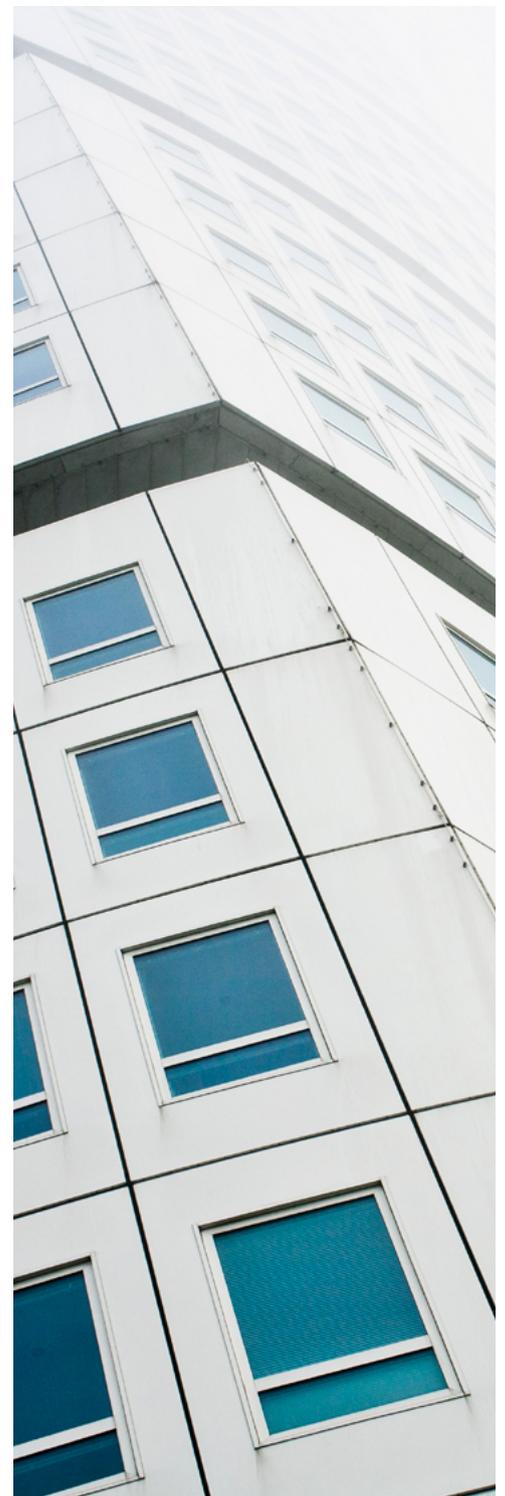
Figure 5: Amount of clusters and service types in the system as function of log of the run time of the optimization

CONCLUSION AND FUTURE WORK

We believe that this research subject has great potential for future work and improvements. As deployments are moving towards thousands of clusters on edge environments, we would like to enable linear degradation in the performance of the algorithm to address the issue of scale. In **Figure 5**, we show that the performance degradation of the algorithm is not linear. For example, when a deployment has 1,000 clusters or more, it could take minutes for the algorithm to converge.

Another improvement we want to introduce into our system is better utilizing the autoscale mechanism

implemented by Kubernetes (Horizontal pod autoscaler) and the different cloud providers (Cluster autoscaler). For instance, we might want to scale a certain service to include more pods behind it in some cases. In this way, the service will increase its capacity and avoid sending requests outside the cluster, which will incur higher latency and traffic penalties. We can also improve the algorithm by allowing access to the full application dependency graph. In that way, we can take into account the whole service chain and make sure to get global optimum, even though this problem was proved to be NP-hard. 



Feature



About the Author

Gordon Haff

is Technology Evangelist at Red Hat, where he works on emerging technology product strategy, writes about tech trends and their business impact, and is a frequent speaker at customer and industry events. His books include *How Open Source Ate Software*, and his podcast, in which he interviews industry experts, is *Innovate @ Open*.

Ops is the new code: Operate First brings open source to operations

Operations are attracting increased attention in the open source community, and the open source ethos is evolving to embrace it.

by *Gordon Haff*

The focus of open source was initially on the code. Over time, however, the health of communities creating that code and associated artifacts such as documentation has also become an open source issue. The approach to governing projects and onboarding contributors hasn't replaced historical concerns such as licensing, but it has assumed a more prominent role.

THE VIRTUOUS CYCLE OF OPEN SOURCE DEVELOPMENT

That we still talk about the open source development model is telling. That language emphasizes developers and other participants in that virtuous cycle, such as users and businesses. Thus, fundamentally, the focus is still on the code. But there's also a dawning recognition that just writing code in a vacuum isn't sustainable for most significant projects.

The attention to code is understandable. Even when proprietary software was nearly the only game in town, vendors were focused on delivering packaged bits to users with maybe some consulting on the side to get it running. Open source software freed users from a

vendor's proprietary IP and let them harness the innovation in a community extending beyond a single company. But it didn't really change the software delivery model. Users were still mostly obliged to operate the software by themselves.

THE RISE OF OPERATIONS: OPERATE FIRST

This model is changing in the modern era. Operations are becoming as important as, and sometimes more important than, code. Software-as-a-Service and public cloud providers have increasingly offloaded the operational burden of software from users. This is a challenge for open source software. While the open source development model is powerful, the value of software lies in operationalizing it so that a user can be productive with it.

One approach to dealing with this challenge is to bring something akin to the open source development model to operations. Enter [Operate First](#).

The term Operate First comes from an open source development model best practice, Upstream First. With Upstream First, the goal is to get every line of code into an upstream project

before it ships as a product. This keeps the community project and downstream product closely aligned and reduces the effort of maintaining divergent code trees. An Upstream First approach recognizes that the value of open source lies not so much in the ability to view source code but in fully embracing an open approach to creating software.

You can think of Operate First as a concept, philosophy, and vision to improve open source software through open sourcing operations. In an Operate First environment, open source code is tested and proven under real workloads running at scale as they would in production. This creates a feedback loop for developers seeking to improve code operationally. Operate First and associated initiatives also aim to document how production deployments are architected and deployed. In addition to documenting best processes and practices, the Operate First project will have an Infrastructure-as-Code repository.

WHAT DOES OPERATE FIRST LOOK LIKE?

Concretely, Operate First is a project to define, build, and improve the open source hybrid cloud through learning and developing code and practices in an open production community cloud. By incorporating operational experience into open source software development, Operate First extends development to include operating, testing, and proving code in a production environment—and simplifying the deployment of that code. It builds on and complements

Operate First serves these goals by encouraging and enabling software design that builds in operational capabilities while keeping the person who needs to operate the software in mind.

a variety of nascent and ongoing projects in the cloud space.

Operate First started as a segment of the [Mass Open Cloud \(MOC\)](#) called the zero cluster, a production cloud set up to host projects and developers seeking to operate first. Announced in 2014, the MOC is a production public cloud based on the model of an Open Cloud Exchange (OCX). In this model, many stakeholders, rather than just a single provider, participate in implementing and operating the cloud.

In addition to the MOC, Operate First is closely associated with various overlapping initiatives, including [OpenInfra Labs](#) (under the Open Infrastructure Foundation) and the [Red Hat Collaboratory at Boston University](#).

OpenInfra Labs hosts the [Telemetry Working Group](#), one



of the working groups included under the Operate First umbrella. Observability of infrastructure has become an increasingly hot topic given the challenge of reliably operating distributed systems such as those in Kubernetes environments. The term can cover a lot of ground, but a typical definition of observability spans metrics, tracing, and logging. Monitoring is often considered something distinct, but it's also at least closely related. A key part of observability is the automatic collection and transmission of data about the system. In other words, telemetry. Telemetry is, therefore, an integral component of Operate First.

THE OPERATE FIRST COMMUNITY

The development of a community around Operate First is still in its early stages. A primary goal of that development is recognizing that there are many constituencies with disparate concerns and motivations. Operate First founders want to engage with them in a manner and through a path that those constituencies prefer.

To start this process, community leaders conducted a series of interviews with a variety of different stakeholders: developers, quality engineering (QE), site reliability engineers (SRE), traditional system admins, data scientists, and others. The objective here was two-fold. First, it was important to understand, for each role, their most pressing day-to-day concerns, what motivated them, how they measured success, and what would make Operate First of interest to them. Second, to keep things simple, identifying and combining roles that largely shared motivations and concerns would make it easier to focus engagement efforts.

THE DIVERSE NEEDS OF OPERATE FIRST PERSONAS

Quality engineers who write testing frameworks and tests have an increasing amount of overlap with more traditional **developers** of applications and other code. Both are motivated by improving customer and internal user experiences, especially when doing so involves solving novel problems. They measure success with metrics such as satisfaction of and adoption by their constituencies as well as productivity and code quality metrics. Operate First serves these goals by encouraging and enabling software design that builds in operational capabilities while keeping the person who needs to operate the software in mind.

From an operational perspective, the focus is shifting away from traditional sysadmin roles that deal mainly with maintaining and upgrading hardware and software infrastructure using tools like scripts and configuration management. While those tasks continue, **site reliability engineers (SREs)** spend a significant amount of time on development tasks such as adding new features, improving scalability, and automating. SREs interact extensively with cloud APIs, whether on premises or in a public cloud. SREs aim to do more with less; the ratio of SREs to the number of managed clusters is one important metric, as is their uptime.

In addition to developer and operations personas, the **data scientists and data engineers** in the OpenDataHub community have also been early adopters of Operate First. OpenDataHub is a blueprint for building an Artificial Intelligence (AI)-as-a-Service platform that integrates a variety of open source machine learning tools, including Kubeflow, Kafka, Seldon, PyTorch, and Jupyter notebooks on the Red Hat® OpenShift® Container Platform.

For these audiences, Operate First provides:

- A cluster to develop and run AI applications
- GitHub organizations to share and collaborate on open source projects
- Custom image pipelines to publish reproducible experiments
- Real production operations data for tackling machine learning problems in AIOps

Furthermore, operating a subset of OpenDataHub at scale creates an opportunity to document best practices, which can, in turn, be fed into Red Hat OpenShift Data Science, the managed cloud service offering based on OpenDataHub. Just as the open source development model forms a virtuous cycle when working as intended, [Operate First can lead to a beneficial circle](#) for operational knowledge and supporting code.

FLEXIBILITY AND FREEDOM

The ultimate goal of Operate First is to free software users from having to make a false choice. It brings the power of the open source development model to operationalizing software. Fully operationalized software is software that maintains the flexibility of open source software that isn't tied to a single cloud provider, while also simplifying and improving the Day Two operations of that software.

It's not simplicity or choice.
It's simplicity *and* choice. 

How can your organization support innovation & future OS leaders?

**Become a Sponsor of the
Open Source Research
Experience Program**

Engage with academic researchers and student talent working on exciting open source projects!

Find out about the benefits of supporting the University of California's OSRE program at <https://bit.ly/3qvMCXK>

Or scan here for more details



Feature



About the Author

Jennifer Wood

is the manager of the Centre for Doctoral Training in Cloud Computing for Big Data at Newcastle University. She is also the University Liaison Manager for Newcastle University's partnership with the Alan Turing Institute in London. She has been working in data science for over six years and has a master's degree in Culture and Difference from Durham University.

Where will we find the data scientists?

Universities play a primary role in developing data skills, but traditional education alone can't close the skills gap fast enough.

by Jennifer Wood

The mismatch between the widespread need for strong data skills and the current workforce is an obstacle for nearly every sector of the economy, which means no single sector can solve it. Collaborative partnerships among universities, businesses, and government agencies have the potential to make the changes needed to bring data science into the mainstream. In this article, I'll describe how the data science program at Newcastle University has worked with government and industry support to find multiple solutions for the data skills gap.

ADDRESSING THE DATA SCIENCE SHORTAGE

Vast amounts of data are now collected across all areas of academia, research, industry, and the public sector. The challenge has shifted from collecting data to extracting value from it.

Everyone needs certain skills to succeed in a data-rich environment, starting with basic data literacy, and the shortage of these skills threatens the ability of businesses and even countries to drive innovation to create economic growth. [The World Economic Forum 2020](#) report found that only 61% of the working-age population in the United Kingdom has digital skills. The UK's skills gap runs across the entire learning journey, from early years to senior positions in academia and industry.

The COVID-19 pandemic highlighted this data skills gap by revealing our reliance on technology. While this reliance created problems for workers who lacked needed skills, it also created the impetus to accelerate digital transformations and demonstrated how technology can lead to success. It's not controversial to say that developing data skills is fundamental to a thriving economy. Industry can't close this gap alone. Universities have a critical role to play, in partnership with both industry and government. The Newcastle Data community (part of Newcastle University) was formed in part to respond to these challenges. It focuses on three main goals:

1. Using data to transform research across the university
2. Training the next generation of leaders in data science
3. Exchanging data expertise with those outside the university, for the benefit of society, the economy, and the university's own research and teaching

The Newcastle Data community coordinates work across research, teaching, and engagement to create a virtuous circle in which research and

work with external organizations—from startups to government and public sector organizations to FTSE 500 companies—keep our teaching up to date. These connections generate a pipeline of talent that feeds both academic research and external organizations.

A COHORT MODEL HELPS CREATE A TALENT PIPELINE

PhD research can be a solitary experience—quite different from how real-world practice usually works. The Centre for Doctoral Training (CDT) model established in the United Kingdom was explicitly designed to encourage collaborative, interdisciplinary research that addresses global problems, supported by both government and industry funding. The Centre for Doctoral Training (CDT) in Cloud Computing for Big Data at Newcastle works with partners to realize this goal as well. Rather than funding single studentships, the CDT funds between eight to twelve students per year and trains them as a cohort. Funding initially allowed us to recruit five cohorts, but thanks to generous industrial sponsors, we've been able to recruit eight cohorts. We now have a dynamic community for both students and staff that prevents learning from being a passive and isolated experience.

Solving real problems with data requires people with deep knowledge and skills in both computing science and statistics. Practical experience in cloud computing and handling real data sets is also required. Traditional single-subject programs don't provide

graduates with this combination of skills. The great advantage of CDTs is that they encourage interdisciplinary work, which gives us the freedom to design our CDT to target these areas.

Students gain a great deal academically from this collaborative approach. Working on real-world problems gives an authentic purpose for their research and opens it to a broader audience. Because they are co-supervised by colleagues in computing, maths, statistics, and other disciplines, they get expertise and perspective from fields other than their own. A cross-disciplinary view enhances students' research in many ways, enabling creativity and innovation but also helping them understand the limits of their knowledge.

Throughout their teaching and research, students share a dedicated office space that enables them to take advantage of the range of backgrounds, knowledge, and skills from across all cohorts. Students also undertake several modules that strongly emphasize group projects, working on a current problem for one of our industry partners. This program design paves the way for future collaborations and working partnerships, and it helps the students develop robust conflict management and relationship-building skills.

IT'S NOT JUST ABOUT TECHNICAL SKILLS

The CDT also provides the flexibility for universities to help students build

business leadership and entrepreneurial skills—again, something that is not common in standard graduate research programs. These abilities are much needed. The recent [Quantifying the UK Data Skills Gap](#) report identified that around a quarter of businesses said graduates who work with data need to develop their leadership and communication skills.

We designed our CDT with this in mind. When we began in 2014, we wanted to develop a program that would produce future leaders in data analytics. This requires not just technical knowledge but also the ability to generate and pursue new business opportunities, either through start-ups or in existing companies.

The challenge has shifted
from collecting data to
extracting value from it.

Universities are seeing a demand for graduates with a skill set that includes core professional skills such as critical analysis, communication, and creativity. We developed a successful collaboration between Newcastle University, the [National Innovation Centre for Data](#) (NICD), and AkzoNobel (a Dutch multinational company) on a data-driven innovation module. The module equips students with commercial awareness around the use of data and AI through a ten-day incubator where students

are immersed in an industry setting, collaborating on real business problems. Using the business model canvas (an entrepreneurial technique) as the foundation for problem solving, the students work through several iterations of a solution to offer fresh, creative perspectives on a traditional company.

At the end of the two weeks, students pitched their solutions to stakeholders within AkzoNobel. According to Mo Chowdhury, AkzoNobel Innovation Incubator Project Lead, “The dedicated and sprint-like mentality provided us with business models that would have taken much longer to produce. Each idea was truly transformative.” Given the potential in using entrepreneurial techniques to foster innovation, this program has also been adapted to suit undergraduate students from non-technical programs.

WIDENING PARTICIPATION

The skills gap creates excellent opportunities to diversify and support students from groups historically underrepresented in tech. [Over 70% of the 1.5 million roles at risk of automation](#)—including artificial intelligence and future technologies—are held by women. The inequitable impact of COVID-19 on women and Black, Asian, and minority ethnic (BAME) communities has slowed progress on diversity and inclusion in all sectors. Meanwhile, progression into postgraduate training is as low as 11% for Black students and 8.4% for disabled students.

To strive toward correcting this, Newcastle used funding from the UK



Students and industry representatives meet in the new Catalyst space

government’s Department for Digital, Culture, Media, and Sport (DCMS) and the Office for AI (via the Office for Students) for a project to widen participation in data science and AI. The project includes forty-five Master of Science scholarships for historically underrepresented groups in the field, with a focus on female, Black, and registered disabled students, students from POLAR Q1 and Q2 (a UK measure of educational participation by locale), care leavers (i.e., a person who spent time in foster or residential care), estranged students, Gypsy/Roma/traveler students, refugees, children from military families, veterans, and partners of military personnel. Dr. Matt Forshaw, a senior lecturer in data science at Newcastle, used what we’ve learned from combining skills in computing and statistics and developed the suite of MSc courses in data science.

UNLOCKING THE POTENTIAL OF THE CURRENT WORKFORCE

Universities play a vital role in creating a talent pipeline of graduates, but with [80% of 2030’s workforce already in](#)

[employment](#), we need different solutions for short-term change. Reskilling and upskilling the existing workforce is imperative. One difficulty is the widespread belief that there is only one pathway to working in fields like data science. A 2020 [Europe-wide YouGov survey](#) (commissioned by Red Hat) highlighted the misconception that only those with data-related qualifications can pursue a career in data or tech.

Bridging the skills gap among people already employed in other fields requires industry, academia, and government to build new pathways and make them achievable. In the United Kingdom, for example, an Apprenticeship Levy creates funds to support employer-based apprenticeships that teach employees of any age and career stage new skills, from data literacy to data analysis and AI/machine learning.

One of the successful strategies developed at Newcastle is supporting organizations by helping their



The Newcastle Helix serves as a hub for collaborative public-private research in data science

existing workforce gain the skills and knowledge that has traditionally been the purview of academic experts. The NICD is capitalizing on the new Catalyst facility, located in the Newcastle Helix, which was specifically designed to bring together researchers and businesses to share the wealth of skills and knowledge currently locked within universities.

A technical team from the NICD, including several CDT PhD and data science MSc graduates, works alongside organizations facing data science challenges, addressing specific needs or data problems. Unlike a traditional consultancy, the NICD technical team works both to find tangible data-driven solutions for clients and to upskill

employees of clients' organizations. As a result, the organization's workforce will be able to tackle the next data project themselves.

THE FUTURE OF DATA SKILLS TRAINING

A great deal remains to be done to solve the skills gap in the United Kingdom and elsewhere, but partnerships like those created via Newcastle's data science program will play an essential role in meeting the need for data skills at all levels. We've seen the benefit of a cohort model and, with the development of the Europe RIG, we expect our relationships with partners like Red Hat and other businesses to grow.

Mark Little, who is both Vice President of Middleware Engineering at Red Hat

and a visiting professor at Newcastle University, leads the Research Centre at Newcastle. Recently named a [Fellow of the Royal Academy of Engineering](#), Professor Little points to a history of success achieved by the joined forces of academic, industry, and government: "Red Hat and Newcastle University have worked together for many years with a track record of successes including five-star rated EU projects, PhDs, upstream open source projects that have been adopted by various companies and academic institutions, and creating new leaders in R&D for Red Hat and other organizations. As hybrid cloud, edge/IoT, and data science research opportunities continue to grow, it is these kinds of successes which we should build upon and strengthen our partnership." 

Column



About the Author

Matej Hrušovský

has been with Red Hat for more than eight years, six of which have been spent managing the university program in EMEA. Aside from attracting new talent mainly from universities and schools, the core of Matej's job is to find and put the right people from Red Hat and academia in the same room together.

Setting the standard for PhD support in the Czech Republic

Collaboration between industry and academia has become a popular talking point in tech, but that wasn't always the case.

by *Matej Hrušovský*

When Red Hat Czech took the first steps towards supporting PhD students directly, this type of funding and mentorship was not common. The initiative began in 2015, in cooperation with the Faculty of Informatics at Masaryk University in Brno, where significant support continues. The growth and continued success of the partnership between the two institutions, founded on a program aimed at undergraduate and graduate students, made direct doctoral support a natural next step. Starting small and working iteratively, the program has inspired a number of other partners to follow in Red Hat's footsteps.

These are some of the most essential factors for the initiative discovered so far:

First: Supporting individual students rather than an entire institution. This direct method of support makes it possible to select projects that resonate with an open source vision.

Second: Minimizing control over the PhD student's research. Red Hat provides financial support but does not contract the student for their job; it merely enables them to do better research. And while the company does have a right to comment on the

student's progress and direction a couple of times a year, it does so from an advisory position, shaping the direction, not prescribing specific work.

Third: Providing a community network and the know-how of working engineers. Accomplishing this means finding a cooperating engineering team or even a "co-advisor" among employees. Finding them, identifying a topic interesting to both parties, and maintaining working communication is the most challenging part to get right, and best practices are still evolving.

As in any project of this kind, some efforts bring more results than others. Not all of the early academia-industry collaborations yielded tangible results, but in all cases getting industry support improved conditions for those students and gave them the freedom to work in the open. This was an important step forward in itself, but the primary achievement of the program was yet to appear.

One aim of offering PhD support—also a primary goal of Red Hat Research itself— was to improve the educational ecosystem overall by creating a new standard for cooperation between the industry and academia. This wasn't a program

meant to build relationships only within Red Hat. Other companies were not only welcome to copy the program—replicating the program elsewhere was the desired outcome. The more companies that make similar efforts, the easier it is for such cooperation to become the established approach.

This is exactly what happened. Universities began to adopt this model of partnership with industry. Tech companies began developing similar best practices for forging connections in academia. Very recently, this trend has continued with PhD support as well. This new model of PhD support, which took Red Hat Czech years to polish, is now being adopted by other companies in the region.

Since our partnership first began, Masaryk students, with Red Hat support, have made advancements in the research of usable security design, automating software verification, AI/ML for adaptive learning, and several other areas. These projects are just two examples of noteworthy accomplishments:

Martin Ukrop: Investigating usability of security APIs, focusing on TLS certificate manipulation

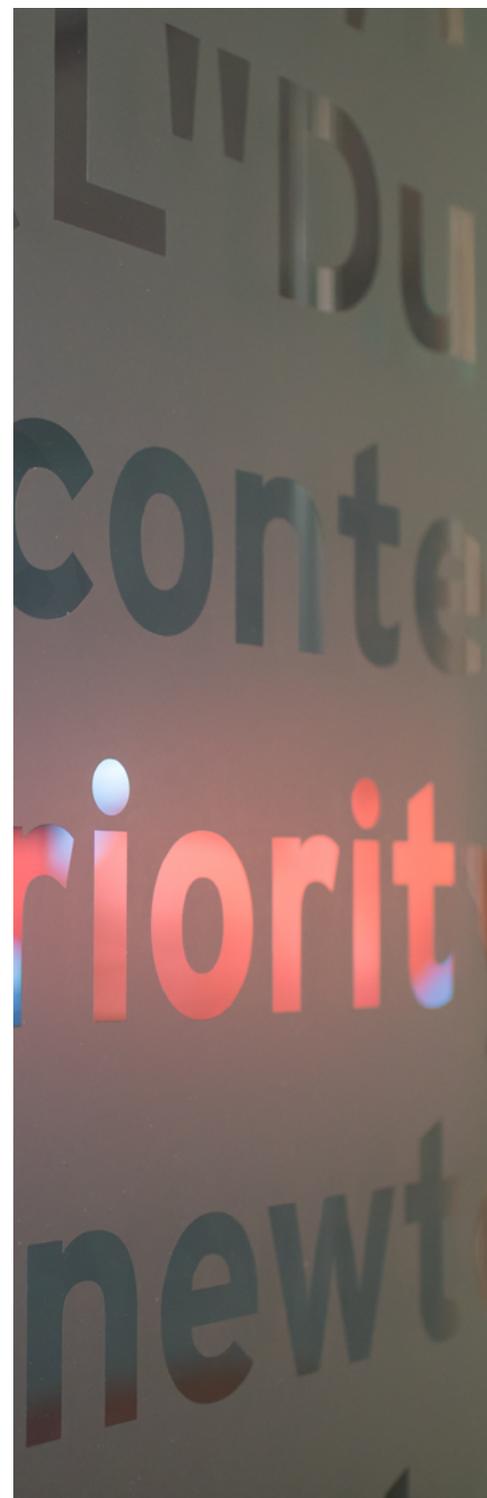
Through his relationship with Red Hat, Martin set up a research booth at DevConf.cz, an annual technical conference. He was thus able to interview many developers year-over-year and evaluate their first-hand experience with security software. Furthermore, the know-how and connections of

Red Hat engineers helped Martin get changes suggested by his research directly upstream into the OpenSSL codebase. Martin's article "[Don't blame the developers: Making security usable for IT professionals](#)," published in issue 2:2 (August 2020) of this magazine, is just one of the results of this successful cooperation.

Agáta Kružíková: Investigating authentication in open source repositories

Thanks to support from Red Hat, Agáta was able to investigate the behavior of developers in real-world scenarios instead of using students in a lab setting. As a result, the research was much more ecologically valid and generalizable. Agáta's work allowed both researchers and Red Hat to better understand security risks related to open source dependencies. Like Martin, Agáta published an article about her work in an earlier issue of this magazine: "[User authentication for open source developers: What do they use?](#)" (3:2, August 2021).

It is exciting to see how far simple cooperation can go when backed with an honest vision, open communication, and mutual interest. The PhD support program's path to success also highlights the importance of courage when a new approach needs discovering. It may not work on the first try, perhaps not even on the second, but this method supports a core open source value: the best ideas always win, no matter where they come from. 



Project Updates

Research project updates

Each quarter, *Red Hat Research Quarterly* highlights new and ongoing research collaborations from around the world. This quarter we highlight collaborative projects in the greater Boston area of the United States.

There are many more active projects, so be sure to check research.redhat.com listings. You also can join live Research Interest Group (RIG) presentations each month to discuss new project proposals and review the latest results from other research collaborations. Subscribe to the [US research mailing list](#) to stay current on the interest group meetings.

PROJECT: Unikernel Linux

ACADEMIC INVESTIGATORS:
Orran Krieger, Renato Mancuso, Ali Raza, Thomas Unger, and Parul Sohal (Boston University)

RED HAT INVESTIGATORS:
Richard Jones, Larry Woodman, and Ulrich Drepper

The Red Hat Collaboratory's Unikernel Linux (UKL) research project expanded its investigation of adapting unikernels to Linux. In situations where the performance of one (trusted) application is critical, the team is experimenting with co-optimizing the application and the Linux kernel. With just a simple recompilation, dozens of

applications tested thus far have shown modest performance gains. For example, Redis's 99th tail latency is improved by 10%, and throughput is improved by 21%. With more effort, expert developers can perform deep optimizations that call internal kernel functionality and employ techniques that enable link time optimizations across the application/kernel boundary. With only ten lines of code changed, Redis's tail latency is improved by 23%, and its throughput is improved by 33% using the unikernel.

While UKL can allow optimization for only one process at a time, standard scripts can be used to launch that process, and other processes can run alongside it, enabling the use of standard Linux user-level tools and infrastructure. UKL supports both virtualized and bare metal x86-64 systems. Although the team has not yet done extensive testing, UKL should support all Linux devices and accelerators. Researchers are reviewing UKL's modest (<1500 LOC) changes to the Linux kernel with collaborators and plan to submit a patch to the Linux kernel mailing list (LKML) early in 2022. The UKL code is available at github.com/unikernelLinux/linux.

PROJECT: High-performance certified trust for cloud-scale enclaves

ACADEMIC INVESTIGATORS:
Zhong Shao, Richard Habeeb, and Hao Chen (Yale University)

RED HAT INVESTIGATOR:
Bandan Das

This project, which investigates the feasibility of building trusted, formally verified computing enclaves for the ARM and x86 platforms, has implemented more of its formal language for specifying and composing enclave layers. Remote attestation libraries and a lightweight version of the ROS-like middleware (called ThinROS) are now running on top of the CertiKOS hypervisor kernel. The ARM platform boots CertiKOS (with real-time support) in the secure world (under TrustZone) and standard Linux in the normal world. A newly implemented security monitor allows normal-world Linux to be treated as if it were a VM process running on top of the secure mode CertiKOS. The x86 platform work explores ways to support VM-based enclaves and I/O device passthrough,

based on a lightweight KVM-based Type 1.5 hypervisor. In related work at Yale, the group investigated how to formally verify the isolation property of a distributed memory manager for a large-scale datacenter with a disaggregated architecture.

PROJECT: Fuzzing device emulation in QEMU

ACADEMIC INVESTIGATORS:

Manuel Egele and Alexander Bulekov (Boston University)

RED HAT INVESTIGATORS:

Bandan Das and Stefan Hajnoczi

The QEMU fuzzing team will share new results at the 31st Usenix Security Symposium in August 2022. This work, conducted as part of the QEMU project (github.com/qemu/qemu), presents the Morphuzz hypervisor fuzzer, which reshapes virtual device input space to generically fuzz complex I/O protocols. Virtual devices are a critical interface enabling cloud environments. This work identifies means to insert a fuzzing framework that exercises initialization, teardown, control, and data flow paths without the need for expert knowledge of device internals. Morphuzz is now continuously fuzzing new changes to QEMU code. It provides 81% coverage over 28 virtual device implementations and has identified more than 66 unique and reproducible crashing bugs, all of which were reported to QEMU developers.



The team's current work applies techniques similar to Morphuzz to simplify the problem of kernel system-call fuzzing. A snapshot-based kernel fuzzer can achieve high coverage over complex interfaces, such as KVM, without the need for detailed grammars or descriptions. By reducing the manual effort required for fuzzing, developers can allocate more resources toward fixing bugs and implementing safer interfaces, ensuring a safer environment for all cloud users. A [prepublication PDF](#) is available.

PROJECT: An abstraction for diagnosing performance problems in distributed applications

ACADEMIC INVESTIGATORS:

Raja R. Sambasivan (Tufts) and Mark Crovella (Boston University)

RED HAT INVESTIGATOR:

Juraci Kröhling

A new joint project including Tufts, Boston University, and Red Hat will build and evaluate the efficacy of tools to mine motifs from distributed traces and diagnosis tools that operate on request-workflow traces.

Diagnosing performance problems in distributed applications continues to grow more challenging. One cause is the mismatch between the powerful abstractions developers use to build increasingly complex distributed applications and the simple ones engineers have available to diagnose problems in them. This project proposes a novel abstraction for performance diagnosis, called the Workflow Motif, which describes frequently recurring processing actions. Raja Sambasivan, a leader of the Open Telemetry Working Group, hopes to apply these tools to telemetry data for large datacenter environments collected and shared through that working group. The researchers plan to publish code, documentation, and datasets used for their work with an open source license and contribute to the relevant communities fixes for any bugs they find and diagnose using the new tools. 

Contact academic@redhat.com for more information on any project described here, or visit the [Red Hat Research Project Directory](#) on our website.



AI ON INTEL®



**NOW BUILD THE AI YOU WANT
ON THE CPU YOU KNOW.**

Learn more at ai.intel.com