May 2019 – Volume 1: Issue 1

RESEARCH QUARTERLY

VOLUME 1:1

Bringing great research ideas into open source communities

IN THIS ISSUE:

- **Daniel Gruss – What can we do to improve security and resistance to the Spectres and Meltdowns of the future?**

Also in this issue:

- **IoT: Building an open test automation framework**
- **The ROSE project: Using open source to bring students together**

Red Hat

RESEARCH QUARTERLY

VOLUME 1:1

Red Hat

Table of Contents

About the Director: Hugh Brock is the Research Director for Red Hat, coordinating Red Hat's research 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.

FROM THE DIRECTOR

How often in life do you get the chance to start something completely new? Once a decade? Once a year? Well, I've hit the jackpot over the last six months—not only has Red Hat® asked me to launch Red Hat Research, I have been able to bring you the first issue of the brand new *Red Hat Research Quarterly*. Red Hat Research is devoted to bringing great research ideas into open source communities, and RHRQ is our way to tell you about them.

As you'll see from a quick flip through the magazine, our efforts at Red Hat Research range from projects in security and big data, to student outreach in Israel, to the groundbreaking Red Hat-Boston University partnership called the Red Hat Collaboratory. Our aim throughout is to use our engineers' expertise in open source projects to focus research at universities, governments, and in industry in an upstream direction. We've found willing partners at universities all over the world, researchers who are anxious to step beyond the traditional academic paper-writing arena and accept our help in bringing their ideas into practical use in the open.

We have had a fantastic time building Red Hat Research and creating this

first-ever magazine. We hope you enjoy reading it as much as we enjoyed putting it together, and that it will inspire you to suggest other ways we can use the open source development model to improve research and education where you are.

RED HAT RESEARCH: MOVING GREAT IDEAS INTO OPEN SOURCE COMMUNITIES

If research is all about creating and monetizing intellectual property, should a company that holds no IP even bother with it? Turns out there are some really good reasons why we should. Like everything we do at Red Hat, our approach to research is unconventional, and rooted in our tradition of open collaboration. Done right, the results benefit universities, Red Hat, and open source developers and users worldwide.

At first glance, Red Hat seems an unlikely player in the world of advanced research in software. We have neither the size nor the scope of the famed research organizations at companies like Microsoft and Google, and we cannot and will not monetize the results of research through the hoarding and licensing of intellectual property. Nonetheless, we are the standard bearer for the open source development model. As the largest firm by far based on this model, we therefore have a strong interest in the rapid movement of research ideas into open source projects, rather than into the patent vaults of our proprietary competitors.

Advancing this interest, by fostering strong connections with great universities and connecting our engineers with their researchers, is the mission of Red Hat Research. Although the task is daunting considering the size of the renowned industry research groups out there, we bring some unique strengths to bear. First, we are the only sizeable company in the world that can credibly say to a university researcher that we have no interest in locking up any aspect of their work to monetize it—instead, we seek to bring it into communities where it will benefit us as well as the common good. Second, we have great engineers who work in those communities. They offer their deep understanding of not only the technology but also the people and the politics of a community, to help researchers make their work relevant to a project and eventually a part of it.

One great example of this kind of collaboration is a joint project between Red Hat, Boston University, and Boston Children's Hospital called ChRIS (ChRIS Research Integration Service). ChRIS (http://www.bu.edu/rhcollab/projects/radiology/) aims to create a stable platform for medical image processing. It hosts plug-ins that have the potential to save a great deal of time for radiologists and even have real clinical impact. We got involved with ChRIS through Boston University, where we co-teach a course on cloud computing with Dr. Orran Krieger (https://www.bu.edu/eng/profile/orran-krieger/), featured

"If research is all about creating and monetizing intellectual property, should a company that holds no IP even bother with it? Turns out there are some really good reasons why we should."

elsewhere in these pages. ChRIS was one of the featured projects in the course, and we soon found that we could contribute a lot to its development by moving it to the Red Hat® OpenShift® Container Platform and reworking the UI using the PatternFly user interface library. If you are reading this article at the May 2019 Red Hat Summit, don't miss Red Hat Chief Technology Officer Chris Wright's keynote, where he will talk about the latest additions to the project and the major impact we hope it will have on the medical image processing world.

The ChRIS project shows how Red Hat Research works. We found a project with great potential that needed some help and course correction if it was going to find a lasting home upstream and with that, the potential for broader adoption. We engaged with the project founder, Rudolph Pienaar (http://www.childrenshospital.org/research/researchers/p/rudolph-pienaar), helped him attract interest in the community around his project, and led the work on porting the project to a more stable cloud platform and UI. Currently we are adding support to ChRIS for a Secure Multi-Party Computing (MPC) plug-in that will allow radiologists to compare a brain scan to an average derived from real patients at other hospitals. This kind of comparison has been impossible to date because of patient confidentiality requirements. MPC, however, allows the comparison to happen cryptographically without sharing any of the underlying

This kind of comparison has been impossible to date because of patient confidentiality requirements. MPC, however, allows the comparison to happen cryptographically without sharing any of the underlying data.

A Secure Multi-Party Computing (MPC) plug-in that will allow radiologists to compare a brain scan to an average derived from real patients at other hospitals.

data. Of course, we didn't do the research to make MPC possible ourselves; we found another researcher, Mayank Varia (https://www.bu.edu/hic/profile/mayank-varia/) who was working in this area and helped him make his work usable as a ChRIS plug-in.

Another good example of this way of working is the Linux® unikernel project (https://www.bu.edu/rhcollab/projects/linux-unikernels/). A unikernel is an operating system kernel that is compiled together with a single application that will run on it. Because the unikernel is meant for a single application, there is no need for any of the boundaries and permission checking that a standard, multiuser operating system requires. This means unikernels can achieve significant performance gains with certain applications, compared with the same application running on a traditional operating system. The Linux unikernel project is a collaboration between the Boston University operating systems group, BU PhD candidate Ali Raza (https://www.bu.edu/rhcollab/people/ali-raza/), and Red Hat's own Ulrich Drepper (https://akkadia.org/drepper/) and Larry Woodman (https://devconfcz2018.sched.com/speaker/lwoodman). It aims to build the infrastructure and make the changes to Linux that are necessary to use it as a unikernel, in order to make unikernel application deployment practical in the real world. Drepper's early involvement through Red Hat Research was the key to this project getting off the ground. The researchers involved, not being experienced Linux engineers, assumed that a Linux-based unikernel would be quite difficult. Drepper thought otherwise, and with Raza, quickly demonstrated a different approach that makes a Linux unikernel relatively straightforward to pull together.

In the coming months, we will expand the reach and depth of Red Hat Research. We want every Red Hat engineer to be excited to help with a research project and for there to be plenty of opportunities. For this reason, we are building Research Interest Groups in all our major engineering offices to make it easy for Red Hatters and researchers to connect and work together. We will continue to build collaborative projects between our partner universities in Massachusetts and those we work with in Czechia and elsewhere, so we can enable the movement of great ideas across national and cultural borders as well as into upstream communities. We hope that these efforts, together with the community engagement that is at the core of Red Hat culture, will build bonds between Red Hat and the research world that outpace companies ten times our size.

— By Hugh Brock,
Research Director, Red Hat

RED HAT COLLABORATORY AT BOSTON UNIVERSITY ADVANCES RESEARCH, PREPARES STUDENTS FOR SUCCESSFUL CAREERS IN TECHNOLOGY

AUTHORS

– Azer Bestavros, Director, Hariri Institute for Computing, William Fairfield Warren, Distinguished Professor of Computer Science, Boston University

– Orran Krieger, Professor of Engineering, Boston University; Principal Investigator, Mass Open Cloud and Red Hat Collaboratory

From machine learning to virtualization, top research institutions like Boston University (BU) have been a driving force behind many key transformations in technology. Red Hat has played a critical role in driving and productizing open source. Can we find ways to better integrate the innovation of research and the open source community? That's the basic premise of the Red Hat Collaboratory at Boston University. The tight partnership between the two organizations connects faculty and students with industry practitioners working in open source software communities. By sharing experiences, cross-pollinating ideas, and working together on shared projects, engineers and researchers with diverse backgrounds and skill sets can improve open source software and accelerate the pace of innovation. Additionally, through the experiential opportunities enabled by the Collaboratory, students can participate in open source communities, see first-hand how research translates to technology, and obtain the necessary skills for successful careers in technology.

The Collaboratory has its roots in the Massachusetts Open Cloud (MOC), an open source, public-cloud initiative hosted at Boston University. The project is a unique collaborative effort between higher education, government, non-profit entities, and industry partners. The success of the MOC clearly demonstrated how the higher education and industry communities can solve complex technical problems by working closely together. We realized that many advanced research projects could benefit from a close-knit university-corporate relationship—and the Collaboratory was born.

IMPROVING RESEARCH, LEARNING, AND OPPORTUNITIES THROUGH COLLABORATION

While researchers at universities like Boston University embraced open source early on, their primary goal is publishing papers, not participating in open source communities. Open source communities, by contrast, must place the reliability and maintainability of the projects they work on above the innovation that can come from research. In the worst case, these mismatched goals let good research ideas die because they are not implemented with broad adoption in mind.

The Collaboratory's mission is to bridge this gap by formulating techniques

and best practices to help good research ideas find broad adoption in open source projects. We apply these best practices across a range of areas including operating systems, cloud computing services, machine learning and automation, and big data platforms.

One key strategy is to empower researchers and open source developers to work together on common projects (see sidebar). We focus on projects that integrate or enable important research, and require a level of development resources and community involvement not normally found in a pure academic environment. For example, BU researchers and Red Hat developers are working closely together on a Datacenter Scale Data Delivery Network (D3N) project to enable caching for data lakes. The project has resulted in both academic publications and upstream contributions to Ceph (ceph.com). Ceph benefits from new innovations, and the researchers are able to evaluate their advancements using real workloads on a production platform. As another example, last summer we kicked off the Unikernel Linux™ (UKL) project, where OS researchers and key Linux developers work together to enable a new execution model for performance-critical Linux applications that can be easily adopted and extended by the upstream community. This project has already yielded an academic paper presented at HotOS and an article on next.redhat.com—a popular site for developers. In fact, the UKL article was by far the most-read post in 2018.

The Collaboratory is about much more than technology. It is about bringing people together and forging lasting and meaningful relationships. It is about fostering diversity in the workplace and advancing the role of women in technology. And it is about mentoring students, extending public outreach and

The Collaboratory prepares students for careers in technology.

BOSTON UNIVERSITY RED HAT COLLABORATORY PROJECT EXAMPLES

- **Unikernel Linux (UKL)**—Explores how the battle-tested Linux code base can be used to build a unikernel that better supports performance-critical applications.
- **ChRIS Research Integration Service (ChRIS)**—A web-based medical image processing platform that supports various forms of medical imaging such as MRIs.
- **Datacenter-Scale Data Delivery Network (D3N)**—A datacenter caching initiative that improves application performance and reduces storage system and datacenter network utilization.
- **Elastic Secure Infrastructure (ESI)**—An open initiative that enables physical servers to be elastically and securely moved between different services in the datacenter.
- **Tracing**—Three Collaboratory projects focused on enhancing end-to-end tracing as a fundamental model for data-center telemetry.
- **MOC OpenShift servic**e—A highly available, multi-tenant, production OpenShift container service for the MOC community.
- **Open data hub**—An MOC service, which runs on top of OpenShift and OpenStack, and acts an abstraction layer that shields data scientists from the complexities of the underlying infrastructure.
- **North Eastern Storage Exchange (NESE)**—An open data lake, architected to meet the large-scale storage requirements of data-intensive science, engineering, education, and technology initiatives.
- **FPGAs in the cloud**—A project to enable the open source tooling required for users to fully exploit FPGAs to address a wide range of machine learning and scientific problems.

The Collaboratory has its roots in the Massachusetts Open Cloud (MOC), an open source, public-cloud initiative hosted at Boston University.

bringing real-world industry experience to the classrooms and hallways of Boston University.

As the name implies, collaboration is the lifeblood of the program. While traditional corporate-university partnerships are often tactical financial arrangements established for specific research programs, the Collaboratory aims to create an enduring, collaborative relationship between Red Hat and Boston University. For that reason, Red Hat maintains a physical presence at Boston University, where Red Hat engineers coach students, lead professional development sessions, and engage faculty. And Boston University maintains a physical presence at Red Hat, where students participate in internships, fellowships, and research projects.

The close interaction of Boston University faculty and students with Red Hat personnel is critical to the success of the program. We host roundtables, workshops, seminars, and colloquiums to build relationships and exchange ideas. And we provide student learning opportunities through cooperative coding events, internships, career development programs, and other activities.

EXAMPLES INCLUDE:

- **Hackathons**—The TechTogether 2019 Hackathon for Women at BU's Agganis Arena brought over 1,000 high school and college students together to develop solutions to real-world problems like combating fake news, reducing environmental waste, and restoring digital infrastructure in the wake of a disaster.
- **Culture and professional development workshops**—Held throughout the year, these informative sessions give BU students an opportunity to seek career advice from seasoned Red Hat engineers.
- **Red Hat Days**—The BU Center for Career Development hosts an annual networking event that connects BU students with Red Hat engineers and hiring managers, helping many students jumpstart their careers.

FOR MORE INFORMATION

To learn more about the Boston University Red Hat Collaboratory and find out how you can get involved, visit www.bu.edu/rhcollab.

RESEARCH QUARTERLY

VOLUME 1:1

TEAMING UP TO BUILD AN OPEN TEST AUTOMATION FRAMEWORK FOR IOT

The Czech Technical University in Prague and Red Hat are working together to develop a test automation framework and test methodologies for the Internet of Things. The initiative will help quality assurance organizations evaluate large-scale IoT systems using open source tools.

AUTHORS

– Miroslav Bureš, Associate Professor, Czech Technical University in Prague

– Štefan Bunčiak, Manager, Quality Engineering, Red Hat

– Miroslav Jaroš, Senior Quality Engineer, Red Hat

In 2017, researchers from the Czech Technical University in Prague (CTU) teamed up with quality assurance engineers from Red Hat to develop a comprehensive test automation framework and a set of testing practices and methods for the Internet of Things (IoT).

Backed by funding from the Technology Agency of the Czech Republic, the team's mission was to produce a collection of test methodologies, create a test automation framework, and deliver an integrated suite of open tools for effectively assessing large-scale IoT systems. The CTU team members bring extensive experience developing test automation methods and best practices, and advising industry. The Red Hat team members bring hands-on experience testing complex distributed systems.

The University, Red Hat, and the larger community all benefit from this collaboration. CTU student interns gain practical experience to prepare for careers in technology. Red Hat deepens its IoT automation testing expertise. The community gains the test methodologies and toolsets critical for successful IoT deployments.

IOT TESTING METHODOLOGIES PROVIDE GUIDELINES AND BEST PRACTICES

The team produced a complete set of methods and procedures for effectively examining the availability and reliability of full-scale IoT systems. The recommendations include:

- Guidelines for creating an effective test strategy for an IoT project—identifying and prioritizing test cases, developing test techniques, determining success criteria, etc.
- Recommendations for automating tests—determining which tests are good candidates for automation.
- Methods for designing automated tests—modeling an IoT system, developing automated tests to validate component interoperability, and examine system behavior under various loads and conditions.
- System quality characteristics and metrics guidelines—understanding which specific IoT system attributes are important to measure and why.

PATRIOT IOT TEST AUTOMATION FRAMEWORK STREAMLINES QUALITY ASSURANCE EFFORTS

The team also developed an end-to-end test automation framework called PatrIoT

> CTU student interns gain practical experience to prepare for careers in technology.

that helps quality assurance organizations efficiently and thoroughly evaluate large-scale IoT systems using extensible, industry-standard tools. The open framework extends JUnit 5 (a standards-based unit testing framework for Java™) and is easily integrated into popular Continuous Integration and Continuous Delivery (CI/CD) tools.

An interactive smart home model and testbed at Red Hat in Brno, Czech Republic.

PatrIoT features and capabilities include:

- A flexible testbed (test environment) that supports any combination of real and simulated IoT devices, addressing the needs of users with different functional requirements and budget constraints.
- A device emulator that models IoT endpoints (sensors, actuators, etc.) and generates data. The emulator includes a variety of generic devices with distinct characteristics. Users can quickly combine these devices to easily simulate a particular endpoint.
- A container-based network simulator that emulates connections between components. Using the network simulator, users can quickly set up testbeds, incorporate emulated devices, define network topologies, and execute tests.
- A test-case generator that automatically produces test cases for specific scenarios, saving time and effort. The tool automatically generates the test configurations, data, and paths that are most likely to expose system defects.
- An open source-based reporting tool and search engine that captures and stores automated test events, results, and audit trails for review, playback, and analysis.

The entire suite of standards-based tools is part of a unified framework with interoperable components, common documentation, and a uniform API.

The PatrIoT framework is currently in pilot testing with several potential system integrator customers who are developing various IoT solutions such as industrial control systems, smart appliances, connected-vehicle applications, and intelligent sensor networks. The ultimate goal of the program is to release the framework as an upstream open source project. PatrIoT will benefit a variety of constituents including hardware and software vendors developing IoT components, system integrators delivering end-to-end IoT solutions, and end customers introducing and using production IoT systems.

FOR MORE INFORMATION

To learn more about the PatrIoT open source project, visit www.patriot-framework.io.

PROFILE OF AN ENGINEER, MENTOR, TEACHER-IN-RESIDENCE: LANGDON WHITE

As the first Engineer-in-Residence for the Red Hat Collaboratory at Boston University (BU), Red Hat developer Langdon White is excited to be interacting with university students and experiencing the fresh perspective they bring to projects. Langdon sees it not only as an opportunity to share his expertise with students, but also a way to tap into the enthusiasm and creativity these students can bring to the open source community. I recently had the pleasure of interviewing Langdon and learning more about the life of a Red Hat Research Engineer-in Residence.

In addition to being an Engineer-in-Residence at BU, I understand that you are associated with the Metro-Boston Research Interest Group (RIG). Can you tell me what that covers?

Certainly. I actually share leadership of the RIG with Rashid Khan, another Red Hat software engineer. We try to cover universities throughout the Boston area, including Boston University, as well as other engineering-focused schools in the area like the University of Massachusetts at Lowell. Being co-leader of the RIG as well as an Engineer-in-Residence means I wear two hats. My Engineer-in-Residence hat has me working closely with students on projects to help them learn more about real-world applications and to tie them in to resources that can help them achieve their goals. When I put on my RIG leadership hat, I'm looking for ideas that can cross boundaries, are aligned with our company goals, and might be able to get some traction in open source projects.

What are some of the things you do as an Engineer-in-Residence?

One or two days a week I am stationed at Boston University. I hold office hours and students can come in and ask me all sorts of questions. They may have an issue getting something to work and I can help show them how it's been done in commercial applications. Even if it's a question I don't have the specific answer to, I might know someone at Red Hat who can help. A lot of it has to do with student guidance. I talk with students who want to become interns and guide them through that process. I also help teach a class introducing students to software development in industry.

How did you get involved in the program?

I took a position at Red Hat in Boston. Part of the charter of that office is to support the Red Hat Research effort. I was really interested in this. Part of me

LANGDON WHITE

"Recently, I supported TechTogether Boston where 500+ female and nonbinary individuals competed for three days building software for various prizes."

has always wanted to be an academic. Maybe it comes from my parents being academics and growing up with that experience. I also know how quickly things change in Computer Science and I knew it would give me an opportunity to see what the current generation of Computer Science students are exposed to and how they approach software engineering work.

Have there been any particular challenges you've encountered?

Our goal is to try and keep everything as open as possible. Some researchers want to hold on tightly to their intellectual property (IP) until they can leverage it, so sometimes it's difficult to get them to open up.

Are there any specific projects you're working on now?

We have about five different projects going on right now. The first is a collaborative project between Red Hat and Boston University focused on determining the impact of inclusiveness and diversity on open source projects. It just kicked off in January. We're in the development phase, and we have connected the students with upstream resources and started developing software to collect metrics. The second is establishing an open source group at BU. It's a meet up-style group where Red Hat buys the pizza and brings in speakers who help articulate the value of open source and gives students the opportunity to ask questions. The third project connects student developers with our Red Hat desktop team. We've been wanting to do that one for a while to get some fresh ideas. Another project is focused on providing technical "sales" help as part of the BU Spark! program for external nonprofits and city government projects. I am also involved in two related research projects for the Massachusetts Open Cloud to support dynamic allocation of resources using a market model. And finally, we support a number of Hackathons. Recently, I supported TechTogether Boston where 500+ female and nonbinary individuals competed for three days building software for various prizes.

If you were to offer any advice to your colleagues about being an Engineer-in-Residence, what would it be?

Do it! We need a lot more engineers to do this. It's a lot of fun to work with the students—just to see them all starting out once again. In addition, if you have a project that you want to get done, but haven't been able to swing it, you might be able to by tapping into the student resources.

RESEARCH QUARTERLY

VOLUME 1:1

AVOIDING THE SPECTRES AND MELTDOWNS OF THE FUTURE: TAKING A PROACTIVE AND HOLISTIC APPROACH TO CYBERSECURITY

Daniel Gruss is a member of one of the research teams that discovered the Spectre and Meltdown vulnerabilities. In this article, Daniel and Red Hat's Hugh Brock share their views on how the industry can avoid these types of cyberthreats going forward.

AUTHORS

—Daniel Gruss, Assistant Professor, Graz University of Technology

—Hugh Brock, Research Director, Red Hat

Two CPU vulnerabilities have drawn attention to the risks and limitations associated with the ways in which the industry currently approaches cybersecurity. Spectre and Meltdown were discovered by researchers in 2017 and publicly disclosed in 2018. Both exploit techniques used by today's processors to optimize performance. The vulnerabilities cause the processor to transiently operate on data and thereby leak data across contexts.[1]

Had they been discovered and exploited in the wild, Spectre and Meltdown could have had disastrous consequences such as widespread loss of confidential data and use in malware. Fortunately, vendors issued software patches to defend against these threats before they were widely publicized and there have been no reported incidents of bad actors exploiting either of these vulnerabilities in the real world.

The patches circumvented the flaws, but like most workarounds they came with a price. In this case the fixes impaired performance, diminishing the benefits of the CPU optimizations. Ultimately, fundamental hardware and operating system design changes will be required to avoid these types of threats in the future.

The list of related vulnerabilities that have been discovered in the meantime is long and it is constantly getting longer. It is ironic that after Meltdown the general public opinion seemed to be that this was the worst vulnerability to find in this area. Only a few months later, Foreshadow appeared, which went beyond Meltdown—in particular in its effect on fully virtualized cloud systems. And Foreshadow was not the last vulnerability in this series. Today we have more than 21 vulnerabilities on this list. While the research efforts in this area are quickly increasing, a clear taxonomy on these vulnerabilities is still in its infancy.

How did we get into this situation?

Spectre and Meltdown are symptomatic of a larger issue—the IT industry tends

Two CPU vulnerabilities have drawn attention to the risks and limitations associated with the ways in which the industry currently approaches cybersecurity. Spectre and Meltdown were discovered by researchers in 2017 and publicly disclosed in 2018.

[1] *To learn more about these vulnerabilities read "A Systematic Evaluation of Transient Execution Attacks and Defenses," by Canella et. al. at https://arxiv.org/pdf/1811.05441.pdf.*

Going forward we need to put security front and center, and ingrain it into our designs and methodologies. Hardware and software developers and architects must be adequately trained on security threats and best practices, and they must fully consider the security implications of their projects.

to take an ad hoc approach to security. Today, most developers aren't steeped in security. A development team might take basic security considerations into account during the design of a component, but an end-to-end security strategy is typically not a priority. Most teams don't look beyond the security concerns of the individual components they are developing. And security vulnerabilities are almost always dealt with reactively, after a weakness is discovered or exploited.

Going forward we need to put security front and center, and ingrain it into our designs and methodologies. Hardware and software developers and architects must be adequately trained on security threats and best practices, and they must fully consider the security implications of their projects. When necessary, we should bring in independent security experts and testing organizations to examine designs and validate system security on an end-to-end basis.

In many ways, cybersecurity is analogous to healthcare and medicine. Just like human viruses can adapt and become resistant to antibiotics, savvy hackers and cybercriminals can probe systems, identify weaknesses and adapt techniques to inflict damage or steal confidential data. Hence, boundaries that once were safe may become unsafe in the future. A prime example for this is Rowhammer, an effect known by dynamic random access memory (DRAM) manufacturers for decades. Alternatingly activating rows in a DRAM bank at a high frequency can induce bit flips in DRAM cells that were never touched. Today we know that this effect can be exploited by software to obtain root privileges.

The underlying problem is once more the immense complexity of our systems, combined with performance optimizations. The DRAM specifications provide limits for the maximum number of row activations per refresh interval. The CPU, the memory controller, must comply with these numbers. And most of the time, we're very far from these limits, because repeated accesses to the same DRAM locations, which would lead to alternating row activations, are cached. Consequently, it was clear for DRAM manufacturers that caching would prevent row activations at this frequency. CPUs became faster and faster and we are closer to these limits today. But what really changed was the assumed attacker capabilities. Why would the attacker use the cache and not bypass it? It is quite plausible that this attacker capability was simply not known to the right people—again, immense complexity would provide a reason why.

Now, while there was some

documentation about the Rowhammer effect, it was never seen as particularly interesting and was mostly just a side note when discussing how to test DRAM modules before shipping them to the customer. For this reason, it was largely overlooked by security researchers for decades. An unprivileged attacker can use uncached memory or flush memory locations from the cache. This never was a secret. However, it took time for researchers to connect this knowledge with the knowledge of the Rowhammer effect. Once discovered, a multitude of concrete attacks appeared. We can see that while the DRAM limits did not really change, the Rowhammer effect still suddenly became a security threat, because the attacker changed. The attacker became more sophisticated.

Our systems are relentlessly optimized for performance, efficiency, and costs. Our isolated approach to security is no longer acceptable. We must adapt our view on optimizations. We must realize that any boundary we optimize towards today, may well be the security or safety issue of tomorrow.[2]

Going forward the industry must take a holistic approach to security. There must be inherent security assessments in every shortcut and every optimization we take. There must be a vibrant exchange between security researchers and engineers when optimizing towards boundaries, assessing whether these boundaries hold. We must re-evaluate established boundaries. We must build up strategies that work in other areas where we face a more or less intelligent adversary, like bio-organisms did over millions of years. While most human viruses are not fatal, computer viruses often have catastrophic effects.

Hardware and software optimizations and design choices can introduce a variety of vulnerabilities for bad actors to exploit. Just like we take a proactive approach to healthcare with regular preventative checkups, we need to take a proactive approach to security, making it an integral part of our designs. While regularly patching software is becoming more and more the default, we do not have such mechanisms for hardware. Furthermore, we do not perform reassessments before the need for patching arises. In practice, we will never fully eliminate cybersecurity threats. Just like in medicine where we will always have new diseases to cure, in IT we will always have new security threats to contend with. By taking a holistic approach to security, we can strengthen our defenses, reduce exposure, and minimize harm.

[2]*Cost optimizations, for example, may have contributed to the recent Boeing MAX 8 disasters.*

ABOUT THE AUTHORS

Daniel Gruss is an Assistant Professor at Graz University of Technology. He focuses on software-based side-channel attacks that exploit timing differences in hardware and operating systems. He implemented the first remote fault attack running in a website, known as Rowhammer.js. His research team was one of the teams that found the Meltdown and Spectre vulnerabilities published in 2018.

Hugh Brock is the Research Director at Red Hat, coordinating Red Hat's research collaboration with Boston University and other universities worldwide.

RESEARCH
QUARTERLY

VOLUME 1:1

MASSACHUSETTS OPEN CLOUD ADVANCES RESEARCH AND INNOVATION

AUTHORS

—Azer Bestavros, Director, Hariri Institute for Computing, William Fairfield Warren Distinguished Professor of Computer Science, Boston University

—Orran Krieger, Principal Investigator, Mass Open Cloud and Red Hat Collaboratory, and Professor of Electrical and Computer Engineering, Boston University

—Michael Daitzman, Director of Engineering and Product, Mass Open Cloud

As demand for compute services moves to large public clouds, smaller players in the infrastructure space find it harder and harder to compete. The Mass Open Cloud is devoted to the idea that an alternative model—a bazaar to the public cloud's cathedral—is not only possible, but necessary for the continued health of the industry and of open source software.

The Mass Open Cloud (MOC) is a unique collaborative effort between higher education, government, non-profit entities, and industry, conceived to overcome the economic restraints and research barriers imposed by today's commercial cloud offerings. The program provides a comprehensive test environment for investigating and prototyping cloud software, helping universities, application developers, and vendors, accelerate the development of new cloud technologies.

Launched in 2014 by Orran Krieger, Professor of Electrical and Computer Engineering at Boston University, and Peter Desnoyers, Associate Professor at the Khoury College of Computer Sciences at Northeastern University, the MOC takes on the challenges posed by the rapid movement of computing into today's large, opaque public clouds. The MOC is meant to provide a cloud based on an open exchange model, where hardware and software service providers can compete in the open to provide users with the best compute capability at the lowest price.

Public clouds help organizations avoid equipment investments and operations expenses, freeing up staff and budget to focus on strategic initiatives. Without an emphasis on open standards and interoperability, however, today's commercial clouds are still closed systems that use proprietary technology and vendor-specific APIs to limit choice and lock in customers.

Today's commercial public clouds are also inherently opaque. Customers have no visibility into the design, behavior, or operation of the underlying infrastructure. These constraints pose a particular challenge to researchers investigating cloud technology or exploring how to best match infrastructure designs with new applications. With no visibility into or control over commercial infrastructure, independent research teams can't collect the detailed technical data needed to conduct experiments or measure performance effectively. In addition, most research teams lack the contacts and wherewithal to engage commercial cloud providers, gain access

to their intellectual property, or influence their plans.

OPEN CLOUD EXCHANGE MODEL TRANSFORMS ECONOMICS AND FUELS CHOICE

In stark contrast to a closed commercial cloud, the MOC is based on a unique Open Cloud Exchange (OCX) model that allows multiple stakeholders to participate in the implementation and ongoing operation of the cloud. By pooling resources and exploiting economies of scale, OCX participants can contain expenses, eliminate inefficiencies, and accelerate investment returns.

Ultimately, with the OCX model stakeholders will be free to offer resources and services to other tenants on their own terms. A university, for example, can stand up infrastructure to support a specific research project, and sell spare capacity to the community to quickly recoup its investments. Community members can pick and choose the specific applications and services that best meet their particular functional requirements and budget constraints. Application and user data can migrate securely from one project or infrastructure to another.

The MOC serves as a breeding ground for open source innovation, where software is continuously developed, integrated, optimized, and enhanced in a real-world cloud setting. The MOC benefits the academic community and industry. Open source developers gain experience operating their applications and services with real users and real workloads in a production environment. Software vendors and hardware manufacturers work out interoperability issues and validate solutions in a real-world setting, before bringing products to market. Researchers gain full visibility and fine-grained control over a real-world cloud implementation, fueling invention and discovery. In time, the MOC will also serve as a fully supported, large-scale production cloud that gives universities and non-profits access to a comprehensive set of cloud-based resources. Collaborators will be able to work on their projects in one place, without the added expenses of moving data in and out of two distinct clouds.

The MOC's core partners include universities such as Boston University, Harvard University, Northeastern University, Massachusetts Institute of Technology, and the University of Massachusetts; government organizations such as the Massachusetts Technology Collaborative and the United States Air Force; and technology companies like Cisco, Intel, IBM, and Red Hat. Each of these partners provide unique expertise and value in support of the program's primary goals.

The MOC serves as a breeding ground for open source innovation, where software is continuously developed, integrated, optimized, and enhanced in a real-world cloud setting.

POWERED BY OPEN SOURCE INNOVATION

The MOC has been used by thousands of students and researchers over the last four years, supporting numerous courses and research projects. Today it serves approximately 400 community members and supports over 10,000 end users.

The Massachusetts Green High-Performance Computing Center in Holyoke, Massachusetts.

Housed in the Massachusetts Green High-Performance Computing Center, the MOC currently provides approximately 3000 cores of commodity Intel compute capacity and 1.2 PB of storage capacity. In the near future, the MOC will connect directly to the Northeast Storage Exchange—a 20 PB data lake that will ultimately scale to hundreds of petabytes.

Current MOC services include:

- An OpenStack IaaS implementation.
- A Ceph object storage service.
- An OpenShift/Kubernetes PaaS implementation.
- An Elastic Secure Infrastructure service that allows community members to easily stand up and reallocate resources.
- An Open Data Hub machine learning-as-a-service implementation.
- A single sign-on service that lets users securely access MOC services using their existing public or private identity management solutions.

EXTENDING THE VALUE OF OPEN SOURCE

The MOC serves as a proving ground for building a large-scale production cloud based entirely on open source software. Working closely together, researchers and developers from academia and industry, identify and address the technology gaps, integration issues, and operational challenges that are often exposed when individual open source projects are combined to create an end-to-end cloud offering.

MOC developers are extending a number of existing OpenStack projects including:

- Adjutant (a framework to help automate admin and user tasks) to support OpenShift.
- Keystone (an identity service) to improve federation.
- Ironic (a bare metal provisioning service) to support Elastic Secure Infrastructure.

The MOC and its partners are working to deliver all the functionality required to provision, secure, monitor, control, and monetize cloud infrastructure and services, providing uniform onboarding, authentication, authorization, show-back and billing, and telemetry capabilities across the entire technology stack, on an end-to-end basis.

The MOC serves as a reference implementation that can be easily reproduced at other datacenters around the world. Various MOC implementations can be federated to create a global OCX. Ultimately, MOC features and capabilities will be incorporated into upstream projects and into commercial products, giving the world a fully interoperable, end-to-end open source cloud solution that transforms economics, improves choice, and accelerates the pace of innovation.

FOR MORE INFORMATION

To learn more about the Mass Open Cloud visit https://massopen.cloud/

To learn more about the OCX model read "Toward an Open Cloud Marketplace: Vision and First Steps" at http://www.cs.bu.edu/fac/best/res/papers/ic14.pdf.

ABOUT THE AUTHORS

Azer Bestavros is the Founding Director of the Hariri Institute for Computing at Boston University. He is a Professor in the Department of Computer Science in the College of Arts and Sciences and holds affiliated appointments in Electrical and Computer Engineering and in Systems Engineering.

Orran Krieger is the lead on the Massachusetts Open Cloud, Founding Director for the Cloud Computing Initiative (CCI) at Boston University, Fellow of the Hariri Institute for Computing and Computational Science & Engineering, and a Professor of Electrical and Computer Engineering at Boston University.

Michael Daitzman is the Director of Engineering and Product at the Mass Open Cloud. He joined the Mass Open Cloud after working in the private industry as a Head of Product Planning, Head of Technical Evangelism, and Director of Product Development at companies such as Nokia and Vecna.

RESEARCHERS AND INDUSTRY DELIVER OPEN SOURCE TECHNOLOGY STACK FOR BIG DATA APPLICATIONS

AUTHOR

—Dimosthenis Kyriazis, Assistant Professor, University of Piraeus, Technical Coordinator, BigDataStack Consortium

The BigDataStack is an open technology stack for big data applications and operations. The project is based on a unique data-driven architecture that optimizes big data analytics performance, efficiency, and scalability by dynamically adapting compute, storage, and network resources based on service attributes, data flows, and application interdependencies.

Growing numbers of mobile devices and IoT endpoints are generating an ever-increasing velocity, volume, and variety of data. By transforming this raw data into meaningful and actionable insights, enterprises can optimize the way they work, accelerate their decision-making, and improve their financial results.

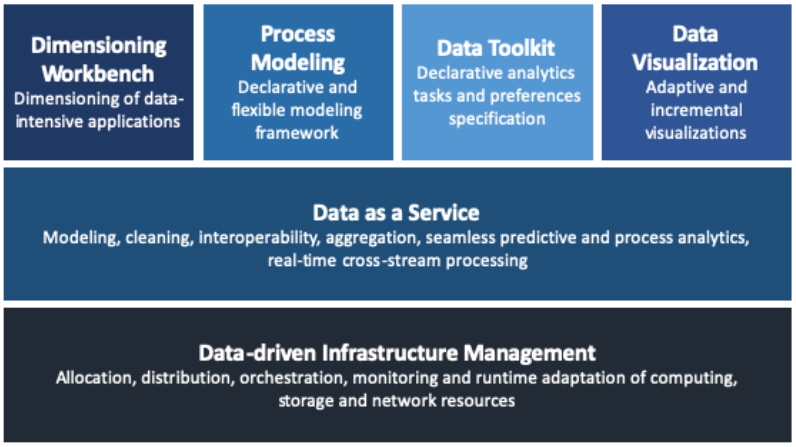

Figure 1. *BigDataStack functional overview.*

Data scientists, business analysts, and application developers require new technologies and tools to analyze these massive datasets in an efficient, cost-effective, and scalable manner. BigDataStack is a data-driven technology stack conceived to address stringent big data analytics performance, agility and scalability requirements.

Funded by the European Commission Horizon 2020 Work Programme, the BigDataStack project will deliver a complete collection of open and interoperable technology building blocks for a variety of big data stakeholders including infrastructure operators, application developers, data providers, data scientists, and data consumers. The BigDataStack project combines business knowledge with academic insights, bringing together engineers, developers, and researchers from the worlds of industry and higher education. Overall, 14 partners are collaborating on the project including Red Hat, IBM, Atos, and the University of Piraeus Research Center.

BigDataStack builds upon and extends the value of a number of upstream open source projects including OpenShift Origin, Apache Spark, and Apache Kafka.

BIGDATASTACK DATA-DRIVEN ARCHITECTURE

BigDataStack delivers an intelligent

and adaptable data-centric framework where compute, storage, and network resources are dynamically allocated based on data flows, service characteristics, and application interdependencies. Unlike with a traditional cloud or datacenter architecture where the underlying infrastructure is focused to meet the computational needs of higher layer applications and services, the BigDataStack features a data-centric architecture that dynamically adapts infrastructure to meet the run-time requirements of data-intensive, latency-sensitive applications. So, for example, instead of indiscriminately provisioning multiple VMs on the same physical server to optimize the performance of a real-time data streaming application, the BigDataStack might intelligently place VMs/containers on geographically distributed physical servers based on data flows/sources locality, data generation rates across the different components of a non-monolithic application or transfer bottlenecks.

Figure 1 provides a high-level functional overview of the BigDataStack.

BigDataStack functional building blocks include:

- **Data-driven infrastructure management**—for dynamically adapting compute, storage, and networking resources in real-time, based on data flows, attributes, and relationships between the services of a composite data-intensive application.
- **Data as a service layer framework**—for modeling, cleansing, manipulating, optimizing, and processing data in motion and data at rest.
- **Dimensioning workbench**—for big data application developers and engineers to specify application and service characteristics and interdependencies, and dimension the application in terms of resource requirements by utilizing AI techniques.
- **High-level process modeling framework**—for business analysts to define and execute business process analytics and optimize business workflows by obtaining recommendations from underlying process mining and optimization algorithms.
- **Data toolkit**—for data scientists to instantiate analytics functions and declare application and data operations preferences and constraints.
- **Data visualization environment**—for users to dynamically view and interpret analytics results, as well as monitor infrastructure and data-related operations.

ABOUT THE AUTHOR

Dimosthenis Kyriazis is an Assistant Professor at the University of Piraeus in Greece and Technical Project Leader for the BigDataStack project. Dimosthenis has participated in several EU and nationally funded projects, leading research for addressing issues related to quality of service provisioning, fault tolerance, workflow management, performance modeling in service-oriented environments, and application domains such as multimedia, post-production, virtual reality, finance, and e-health.

RED HAT DELIVERS DATA-DRIVEN INFRASTRUCTURE

Red Hat contributes engineering talent and software to the BigDataStack initiative. Red Hat engineers are extending a number of upstream open source projects to enable the data-driven infrastructure management layer and operations phase of the BigDataStack program, including:

- **Kubernetes**—an open source system for automating deployment, scaling, and management of containerized applications.
- **OpenStack**—an open source software stack for building private and public clouds.
- **Kuryr**—open source containers networking software for OpenStack.

BIGDATASTACK THEORY OF OPERATION

The BigDataStack theory of operation is characterized in three distinct phases, as depicted in Figure 2.

In the **dimensioning phase**, users specify application and service requirements (related data sources, storage needs, etc.) and data path requirements (services required for data representation, data aggregation, etc.) as illustrated in Figure 3.

In the **deployment phase** optimal implementation patterns are established, based on the results of dimensioning phase, as depicted in Figure 4.

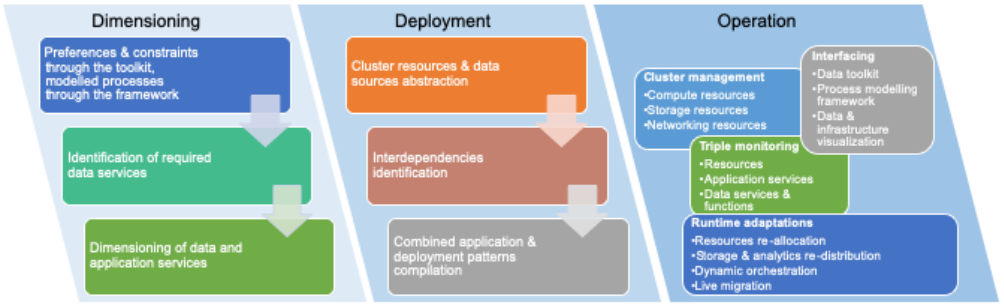

Figure 2. *BigDataStack functional phases.*

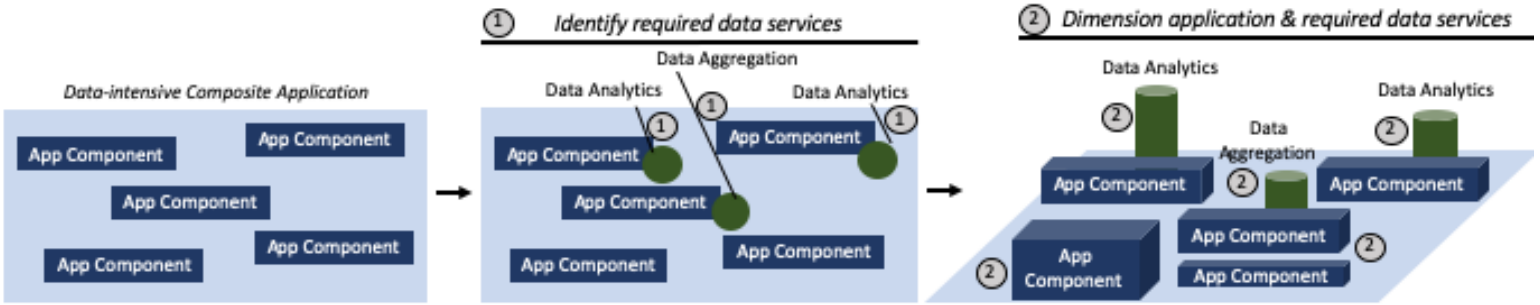

Figure 3. *The dimensioning phase.*

Finally, in the **operation phase** compute, storage and networking resources are dynamically provisioned, monitored, and adapted, based on the results of the deployment phase.

BigDataStack handles the complete data path and life cycle of a big data application, executing analytics in a holistic fashion across multiple data stores and locations, analyzing both data in motion and data at rest.

BIGDATASTACK COMMERCIAL USE CASES

A number of BigDataStack partners are actively developing and evaluating commercial applications for the project. Sample use cases include:

- A **real-time cargo ship management application** to optimize the operation and routing of container ships, enable predictive and preventative maintenance, conserve costs and energy consumption, and improve health, safety, and environmental impact.
- A **connected-consumer application** that gives retailers visibility into customer behavior, preferences, and demographics, allowing businesses to predict future actions and fine-tune sales and marketing programs and shopping experiences based on live data and historical trends.
- **A multichannel financial services application** that gives banks and insurance companies insights into customer tendencies and activities, allowing firms to offer customized services, upsell products, and improve customer satisfaction and loyalty.

These applications help validate BigDataStack performance, scalability, and functionality in real-world scenarios, using real-world datasets. They collect and manage both streaming data and historical data, and will ultimately analyze large datasets of up to 60 TB.

FOR MORE INFORMATION

To learn more about BigDataStack partners, functions, and architectural details visit www.bigdatastack.eu. For detailed technical information download the BigDataStack reference architecture at http://www.bigdatastack.eu/sites/default/files/BigDataStack_D2.4_v1.0.pdf.

The BigDataStack project combines business knowledge with academic insights, bringing together engineers, developers, and researchers from the worlds of industry and higher education.

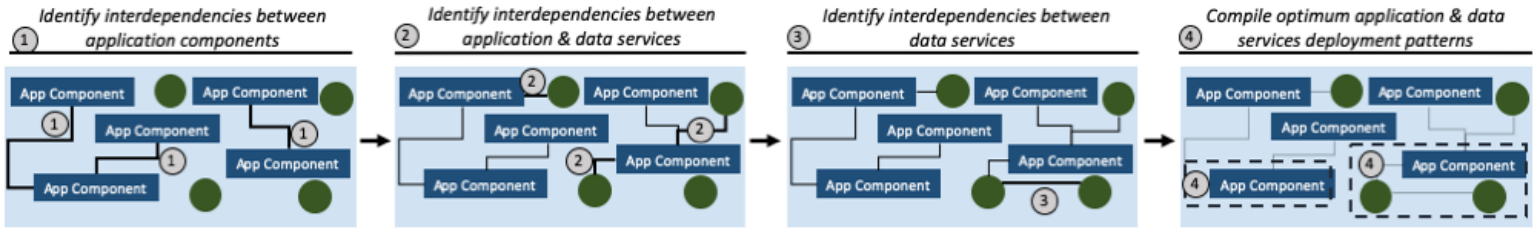

Figure 4. *The deployment phase.*

BUILDING A CULTURE OF COLLABORATION THE OPEN SOURCE WAY

Open organizations like Red Hat value transparency, inclusivity, adaptability, collaboration, and community. Employees come from all over the globe, united in the belief that working in an open way has the power to bring about a better world. One of the best examples of this can be found with the ROSE (Red Hat Open Source Education) project, a cross-community effort that brings students from very different backgrounds and cultures together to learn about the value of open collaboration.

ROSE project students learning about computer programming.

The ROSE project, run out of the Red Hat offices in Israel, offers the opportunity for Israeli students from the Yonatan Middle School in Ra'anana to gather with Arab students from Tira to learn about computer programming using Python and the Linux™ operating system. Started about five years ago, the program is conducted by a group of Red Hat employee volunteers who are passionate about mentoring children, technology, and passing on the values inherent in an open source culture.

When the originators of the program first started, many pointed out the numerous obstacles to success. They faced cultural and language barriers, as well as technical ones. How could they teach a group of students, all who came from very different backgrounds, didn't speak the same language, and hadn't yet been taught programming, how to code in the open in only 14 two-hour sessions?

To address the technical issues, the founders, Miki Kenneth (Director of Software Engineering at Red Hat Israel and the Red Hat Israel site leader) and Eli Masika (a Red Hat Senior Software Engineer), tapped into the open source community for ideas and help. Red Hatter Doron Fediuck came up with an idea for a game—a car race where students could design their own car and race it against the other students' cars. The framework was expanded as an open source project and is available on GitHub. Laptops were supplied so each student was working on an equal footing.

The group tackled the cultural issues in an open source way as well, teaching by example. The mentors came from a variety of backgrounds. There were instructors who could speak Hebrew, Arabic, English, French, and Russian.

No matter how the students wanted to communicate, there was help for them to do so. The fact that the mentoring team was as diverse as it was helped to immediately alleviate any concerns on the part of the students. They saw women and men speaking a variety of languages from differing backgrounds all working together. It made a strong impact on the 9th graders who quickly let down their guard and started to have fun interacting with their teammates. As time went on, the young teenagers soon learned that although they came from different backgrounds, they had a lot in common. They all played on the same game stations, listened to the same music, and had other similar interests.

The instructors ran two sessions of team building exercises that helped break down any remaining barriers. These fun exercises included a marshmallow challenge, an "escape room", and a computer activity where the children teamed up in groups from both schools to take a computer apart to learn how it was built. In the marshmallow challenge, groups were encouraged to build something as high as they could from spaghetti and marshmallows. In the "escape room" the students were challenged to cooperate to solve riddles that would help them exit the room with an important document they needed to bring to the site manager in order to "save Red Hat!"

The focal point of the program, the coding instruction, centered on introducing the students to Linux and Python. The students learned how to use these tools to program their own cars in order to participate in the final exercise, the car race. The students were teamed up in pairs—a student from one school paired with a student from the other—and worked together over the course of the other 12 sessions to design their individual cars. Instructors noted how they could see connections forming between the students. On the day of the big race, one girl joyously celebrated the win of her teammate with a big hug.

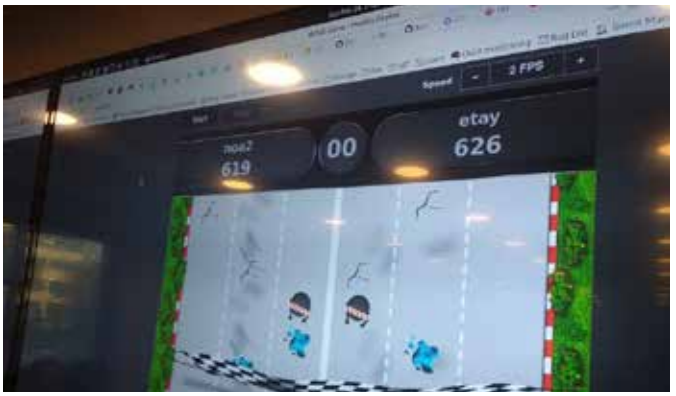

Student-designed car race game.

Students participating in an "escape room" exercise.

While there was some concern that the program might only last a short time, the past couple of years have shown that it is highly valued by the communities involved. The project paused after the first two years to consider if it was worthwhile continuing. Red Hat had been in the driver's seat and the founders weren't sure if they were really making a difference. But after a year-long hiatus, the schools reached out to Red Hat, urging them to continue. As a matter of

The results are rewarding for all involved. The kids learn about open source code while having fun, expanding their knowledge and horizons at the same time. The program gives them skills that can open job opportunities for them in the future.

fact, at the graduation ceremony, not only did the principals of each school attend, but so did the mayors of each town, indicating just how important they viewed this community building effort.

The ROSE project was a huge win for everyone—the kids, the community, and Red Hat. Red Hat Executive Vice President and Chief People Officer Delisa Alexander came to Israel to point out how exposing students to open source will benefit Red Hat in the future. The program emphasizes diversity at its finest, bringing isolated cultures together and laying the groundwork for collaboration. The program has already produced interns for Red Hat and exposed numerous students to Linux who may never have come across it otherwise.

ROSE project participants.

Team members believe that this is a program that can easily be duplicated anywhere in the world. All the components are in place. All that's needed are other teams from within Red Hat to show the same enthusiasm of bringing disparate cultures together and a desire to make it work. It's all about living the Red Hat Way, giving back to the community, and finding ways to collaborate openly. You just need to think about two groups that wouldn't necessarily meet in general and then show them how possible it is simply by emphasizing similarities over differences.

The results are rewarding for all involved. The kids learn about open source code while having fun, expanding their knowledge and horizons at the same time. The program gives them skills that can open job opportunities for them in the future. And Red Hat plants the seeds that will develop into new talent and customers who appreciate the value of open source. It's a win-win all around.

RESEARCH
QUARTERLY

VOLUME 1:1

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

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

RESEARCH QUARTERLY

VOLUME 1:1

In next quarter's issue:

- **Research Day Retrospective. We talk to the professors and attendees at the first-ever Red Hat Research Day to see what ideas they'll be focusing on next**
- **Where it all began: How Red Hat Brno got open source research off the ground in the Czech Republic**
- **Devconf US Preview -- The talks, the ideas, the hula hoops!**

Bringing great research ideas into open source communities

May 2019 – Volume 1: Issue 1

Red Hat