

Red Hat Collaboratory at Boston University Request for Proposals for 2024 Awards

September 7, 2023

Heidi Dempsey, Research Director, US

Hugh Brock, Director

Red Hat Research

Orran Krieger, Professor,

Ari Trachtenberg, Professor

Boston University



Motivation and Track Record

- Red Hat funds Collaboratory projects to:
 - Develop and verify ideas that might not arise on their own in open source
 - Accelerate innovation in systems engineering, particularly in distributed systems, AI for systems, and the edge. The MOC Alliance, NERC, and the MGHPCC make BU very attractive for this kind of work.
 - Get access to live systems with real users in the research community. This presents new research problems and also improves the open-source solutions at Red Hat and upstream.

- Our list of successful projects in the last seven years is good, but we have learned some things:
 - Successful projects have Red Hat engineers involved, but they may have to come from Red Hat Research
 - We need to be very specific about the direction of projects we approve, but not be blind to interesting things that happen along the way
 - Successful projects take advantage of NERC

About Collaboratory Projects

- Part of long-term relationship between Red Hat and Boston University: collaborations with real impact on cloud systems engineering
- Combining academic research with real-world data and experience
- Red Hat open source solutions: linux, cloud, container and Kubernetes, core to edge
- Red Hat Collaboratory Research Areas
 - AI and Machine Learning (esp as relevant to cloud ops and eng)
 - Distributed Systems and Applications (Hybrid Cloud)
 - Composable hardware, software and open hardware tools (Edge)
 - Security and Privacy, including MPC, runtime integrity (Trust)
- Proposals that require Red Hat support
 - Large-scale projects (<\$500,000) with one or more RH engineers committed
 - Small-scale projects (<\$175,000) with one RH Eng at 10% FTE commitment
 - Speculative projects (<\$100,00) designed to initial collaboration

Active Current Red Hat AI Research Areas

- Sanjay Arora, Jason Schlessman - *AI/ML models, training, tuning and applying AI to core systems software and platforms*
https://research.redhat.com/blog/project_member/sanjay-arora/
https://research.redhat.com/blog/project_member/jason-schlessman/
- Michael Clifford - *Red Hat OpenShift AI (formerly RHODS and Open Data Hub)*
<https://www.redhat.com/en/about/press-releases/red-hat-openshift-ai-accelerates-generative-ai-adoption-across-hybrid-cloud>
https://research.redhat.com/blog/project_member/michael-clifford/
- Oindrilla Chatterjee - *machine learning, data science,*
<https://www.bu.edu/academics/cas/courses/cas-cs-549/> <https://github.com/oindrillac>
- Aakanksha Duggal - *data science, software engineering practices*
<https://github.com/aakankshaduggal>
<https://www.bu.edu/academics/cds/courses/cds-ds-219/>
- Lance Galletti - *data science tools and applications, data security and privacy*
<https://www.bu.edu/academics/cas/courses/cas-cs-506/>
https://research.redhat.com/blog/project_member/lance-galletti/
- James Kunstle - *Project Aspen, evaluating health of open source communities*
<https://github.com/oss-aspen> https://research.redhat.com/blog/project_member/237764

AI Examples



Early

Deep-dive into fine-tuning an LLM for documentation

LLMs are both a potentially major workload for RH platforms as well as tools that would be very useful in multiple use-cases across Red Hat.

Status: Started collaboration with Emerging Tech data science team to explore and implement new ML ideas for question-answering tasks from documentation.

Next Milestone:

- For a given set of docs (OpenShift, linux man pages), train/fine-tune an LLM with special tokenizers that might help with CLI commands.

In Flight

Discovering the optimal ordering of compiler passes

Tuning heuristics in compilers has been shown to make a big impact on the performance of binaries. Better methods and models for learning good heuristics provide a competitive advantage for RH's projects like GCC.

Status: Experiments on learning heuristics/orderings for fixed source code complete. Extensions to learn one model that generalizes across source code underway.

Next Milestone:

- Evaluating the effect of using code embeddings from LLMs to learn heuristics that work across a variety of source applications.
- Implementing graph neural networks to learn code embeddings from control flow and data flow graphs.

Nearing Completion

Improving the energy footprint of network-heavy applications

Tuning NIC-level parameters and CPU-level knobs has been shown to have big impacts on performance and energy consumption (and thus, cost) for network-heavy applications. Extending these methods to OpenShift would make the platform more attractive to our customers.

Status: Bare-metal linux experiments for static workloads complete (PhD student graduated in fall 2022). Showed significant improvements in latency and energy.

Next Milestone:

- Experiments using interrupt-level logs for dynamic control of varying workloads on bare-metal linux
- Extension to OpenShift to assess effect of noisier environments on results (Prof Jonathan Appavoo at RH for sabbatical)

Active Current Red Hat Hybrid Cloud and Edge Research Areas

- Lars Kellogg-Stedman, Dylan Stewart, Chris Tate, Tajudeen Salawu - OpenShift, HyperShift, RHODS and Open Data Hub platforms, New England Research Cloud (NERC) engineering <https://github.com/OCP-on-NEERC>
<https://github.com/nerc-project> <https://nerc.mghpcc.org/>
<https://nerc.mghpcc.org/user-guides/>
- Tzu-Mainn Chen, Danni Shi, Surbhi Kanthed, Gagan Kumar - Elastic Secure Infrastructure (hardware resource leasing/management), Curator (container resource usage reporting)
https://research.redhat.com/blog/research_project/elastic-secure-infrastructure/
https://research.redhat.com/blog/research_project/curator/
- Danni Shi, Isaiah Stapleton - Open Education Project build and test infrastructure
https://research.redhat.com/blog/research_project/foundations-in-open-source-education/
- Rashid Khan, Abhimanyu Gosain - LEGO for 6G: a modular open access end-to-end network architecture to enable B5G/6G research, Colosseum testbed tooling
https://research.redhat.com/blog/research_project/enabling-intelligent-in-network-computing-for-cloud-systems/
https://research.redhat.com/blog/project_member/abhimanyu-gosain/

Active Current Red Hat Hybrid Cloud and Edge Research Areas

- Chris Tate, Lars Kellogg-Stedman, Tzu-Mainn Chen, Joydeep Bannerjee, Randy George - Open operations telemetry for research, cluster observability, semantic cloud programming, science applications research with OpenShift, FIWARE
<https://www.bu.edu/rhcollab/events/colloquium-series/>
- Ahmed Sannaullah, Jason Schlessman - CoDes research lab to advance composable hardware projects, prototypes and tools for secure and sustainable edge computing
- Matt Miller, Jeffrey Osier-Mixon, Isaiah Stapleton, Larry Woodman - Fedora and RHEL improvements, RISC-V kernel support, unikernel
- Sally O'Malley - software engineering practices, emerging edge tech, microshift
<https://www.bu.edu/academics/cds/courses/cds-ds-219/> <https://github.com/sallyom>

Active Current Red Hat Security and Privacy Collaboratory Research Areas

- *Victor Malik* - DiffKemp: Automatic analysis of semantic differences in kernel options
https://research.redhat.com/blog/research_project/diffkemp-automatic-analysis-of-semantic-differences-in-kernel-options/
- *Martin Ukrop* - CHESS: Cyber-security Excellence Hub in Estonia and South Moravia
https://research.redhat.com/blog/research_project/chess-cyber-security-excellence-hub-in-estonia-and-south-moravia/
- *Lily Sturman* - Confidential Computing, FHE with FPGA acceleration, Sigstore
<https://research.redhat.com/blog/article/research-perspectives-focus-on-security-privacy-and-cryptography/>
https://research.redhat.com/blog/project_member/lily-sturmann/

Additional Details

- Other Support
 - Most Red Hatters are remote, so reach out now
 - Use the email in the RFP: prop-rhcollab-l@bu.edu
 - Check open source repos related to your interests
 - Contact hdempsey@redhat.com

- More information
 - [Request for Proposals Details](#)
 - [DevConf](#), developer videos (global Red Hat)
 - [Red Hat Research website](#), [people](#), and [project](#) pages