



Security and Privacy Heterogeneous Environment for Reproducible Experimentation

Jelena Mirkovic and Brian Kocoloski (USC/ISI), David Choffnes and Daniel Dubois (Northeastern University)
Geoff Lawler, Christopher Tran, Joe Barnes, Terry Benzel, David Balenson, Srivatsan Ravi, Ganesh Sankaran, Luis Garcia, Alba Regalado

Societal Need

- Our nation depends on correct and reliable functioning of network and computing systems
- Frequency and impact of cybersecurity and privacy attacks are constantly increasing:
 - Solar Winds supply-chain attack, which exposed confidential government data
 - Colonial Pipeline attack, which shut down our major gas pipeline for several days.
 - Ransomware attacks more than tripled
 - DDoS attacks doubled
 - Data breaches increased by 70%
- **Research progress in cybersecurity and privacy is of critical national importance, to ensure safety of U.S. people, infrastructure and data.**

Research Need

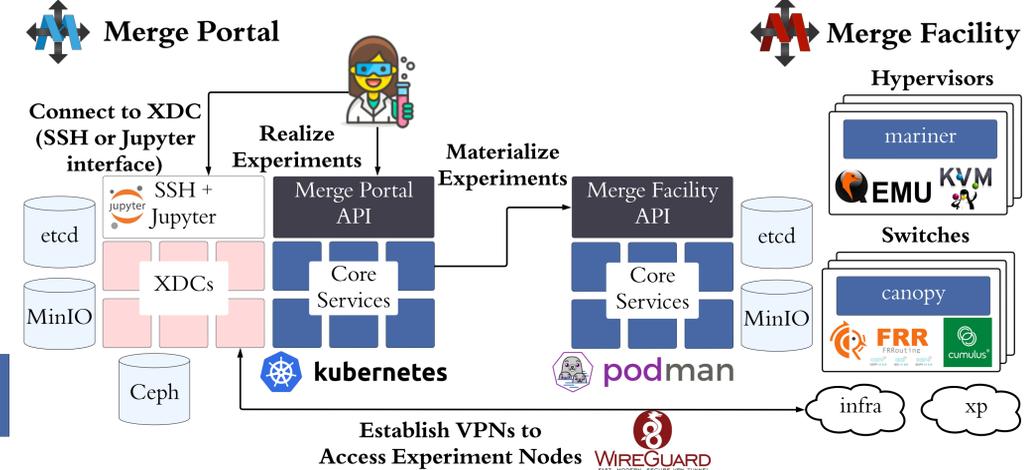
The cybersecurity and privacy research community needs a common, rich, representative research infrastructure, which meets the needs across all members of the community, and facilitates reproducible science

- **Common, rich infrastructure:**
 - Security and privacy issues affect different technologies differently (e.g., different CPU architectures)
 - Some emerging technology can create new vulnerabilities (e.g., IoT)
 - New technologies can be used for defense (e.g., trusted hardware, SDN)
 - Infrastructure must have diverse hardware to meet wide research needs
- **Meet needs across all members of the community:**
 - Experienced and novice users, researchers and students
- **Facilitate reproducible science:**
 - Help researchers create, share, and reuse research artifacts

Merge SW for Research Infrastructure

Microservice Architectures for Modularity and Resilience

The Merge portal and facility codebases use microservice architectures to flexibly integrate homegrown and 3rd party services to implement the Merge APIs



SPHERE Research Infrastructure

Diverse hardware to support diverse research needs (85% of today's publications):

- General and embedded compute nodes with trusted hardware, PLCs and IoT devices, programmable switches and NICs, GPU-equipped nodes

Six user portals supporting:

- Exploratory research (MAN)
- Novice users (GUI)
- Mature research (JUP)
- Use in classes (EDU)
- Use in human user studies (HUM)
- Use for artifact evaluation (AEC)

Libraries of artifacts

- REEs and other artifacts
- Easy reuse on SPHERE

10 GPU-equipped servers
Research supported: security with machine-learning in the loop

600 from DCOMP, 312 new (Intel Atom, Intel Xeon D, ARM Cortex-A57, and NVIDIA Jetson NX Volta GPUs)
Research supported: edge computing security, blockchain security, private computing, trustworthy edge computing, federated learning

48 from DeterLab, 144 new nodes, with Intel TDX, ARM CCA/TrustZone, and AMD SEV
Research supported: application, system and network security, measurement, human user studies, largescale experiments, education, trustworthy computing

8 Tofino switches, 16 Xilinx Virtex-7 NetFPGA development boards (smartNICs)
Research supported: dynamic (programmable) network security, SDN security

500 IoT nodes (a variety of smart home, smart speaker, camera, doorbell, TV, appliance, medical, office, wearable, and miscellaneous devices)
Research supported: IoT security, user privacy

5 Rockwell Automation ControlLogix PLCs, I/O modules
Research supported: critical infrastructure security

Dedicated team of researchers, developers and managers

- Almost 20 years of experience in building and operating testbeds
- Operated the only public cybersecurity testbed - DeterLab
- Built and operated the largest IoT testbed - Mon(IoT)r Lab
- Developed and shared Merge and IoT testbed software
- Developed many tools and datasets for testbed experimentation in cybersecurity and privacy
- Helped others build their own versions of our testbeds at their institutions

Reproducibility support by research infrastructure

- User action logging to alleviate cognitive load
- Help package artifacts on SPHERE
- Support packaging of workflows
- Automatically verify completeness of an artifact and:
 - resilience (does it always run?)
 - consistency of results (does it produce same outcomes within some margin of error?)
 - portability (can an experiment be reused by a different user on a different experiment?)

Flexible security policies:

- Full isolation
- Measurement research
- Software download
- Risky experiments with malware

Transforming Research Community

Need-discovery workshops and surveys

- Presentations and BoFs at major conferences
- Direct engagement with researchers via surveys and interviews
- Discover needs of all community members and adjust SPHERE development to meet them

Help develop standards for artifacts

- Engage wide research community in discussion about artifacts
- Help produce specifications around proper and complete artifact documentation
- Help produce specifications of requirements for quality artifacts

Representative experimentation environments (REEs)

- Used by multiple researchers for a given experimentation task, become a standard for evaluation in a sub-field of cybersecurity and privacy
- Contributed by research community - researchers receive supplemental funding to deploy their high-quality artifacts as REEs on SPHERE

Streamlining artifact evaluation

- Work with artifact evaluation committees (AECs) to have artifacts evaluated on SPHERE
- Artifact authors can submit their artifacts by deploying them on SPHERE
- AECs evaluate on SPHERE, make recommendations for improvement
- Artifacts remain hosted on SPHERE

Broadening participation in computing

- Host 20 minority students per year, involve them in SPHERE development
- Provide research infrastructure to underresourced institutions
- Improve cybersecurity education for everyone via EDU portal and services, hosting of education materials

Visit us at <https://sphere-project.net>

SPHERE is based upon work supported by the National Science Foundation under grant number 2330066. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.