Enabling Reproducibility through the SPHERE Research Infrastructure

Authors: Jelena Mirkovic, Brian Kocoloski, David Balenson
Shepherded by: Rik Farrow

In October 2023, the U.S. National Science Foundation (NSF) funded the Security and Privacy Heterogeneous Environment for Reproducible Experimentation (SPHERE) project via its mid-scale research infrastructure program. SPHERE is a four-year long construction project to build a modern, versatile, and usable common research infrastructure to support cybersecurity and privacy research and education. Led by USC Information Sciences Institute (PIs Jelena Mirkovic and Brian Kocoloski) and Northeastern University (PI David Choffnes), SPHERE aims to transform cybersecurity and privacy research, enabling representative, sophisticated, and reproducible experimentation that allows researchers to build on the work of their peers, thus supercharging scientific progress. The infrastructure is partially complete and already in operation for beta users.

SPHERE also aims to provide usable infrastructure for various classes of users in cybersecurity and privacy areas: both novice and expert researchers, educators and students, investigators running human user studies, and artifact evaluation committees. SPHERE will further enable unprecedented access to hardware and software that is crucial to emerging cybersecurity and privacy fields, such as confidential computing, cyber-physical system security, IoT security and privacy, secure federated learning, etc.

In this article, we describe motivation and need for SPHERE (Section 1), overall architecture, components and services (Section 2), and current status (Section 3). We also explain how using a common research infrastructure helps researchers and educators (Section 4) and enables faster research progress in the entire community. SPHERE is currently open for beta users at https://sphere-testbed.net. Our project page at https://sphere-project.net provides up-to-date information about the project, describes opportunities for collaboration, and outlines plans for the future developments.

1. Motivation and Need

Over the past decade, and especially during the Covid-19 pandemic, our essential functions (e.g., work, school, entertainment, social, financial, critical infrastructure, and governance) moved increasingly online. This sharply increased society's dependence on correct and reliable functioning of network and computing systems, and has led to increases in the frequency and impact of cybersecurity and privacy attacks. Recent years have seen unprecedented and record-breaking attacks, for example the Solar Winds supply-chain attack [1], which exposed confidential government data, and the Colonial Pipeline attack [2], which shut down our major gas pipeline for several days. Ransomware attacks more than tripled [3], DDoS attacks doubled [4], and data breaches increased by 70% [5]. We now live in a world where cybersecurity and privacy are intrinsically intertwined with everything we do, and failures in these domains can have far-reaching monetary and national security impacts, and even jeopardize human lives. Research progress in cybersecurity and privacy is thus of critical national importance, to ensure safety of people, infrastructure and data.

USC Information Sciences Institute ran two workshops in 2022 to learn about community needs around cybersecurity and privacy research: the Cybersecurity Artifacts Workshop [6] and the Cybersecurity Experimentation of the Future 2022 Workshop [7].

Cybersecurity and privacy researchers need common, rich, representative research infrastructure, which meets the needs across all members of the community, and facilitates reproducible science to move from piecemeal, opportunistic research to pursuing integrated, sophisticated, community-encompassing research. We also need a well-educated workforce that is knowledgeable about cyber threats, and that has mastery over practical skills to prevent, detect and recover from cyber attacks.

2. SPHERE Architecture

The SPHERE project is building an advanced research infrastructure (see Figure 1) to support cybersecurity and privacy research and education. Led by the team that built and operated the Deterlab testbed, and supported more than 1,000 research users and more than 20,000 education users over two decades, the SPHERE project has an ambitious goal to meet needs of the broad and diverse research community through modern and diverse hardware, a suite of user portals geared towards different user communities, and a suite of reproducibility services coupled with community-wide efforts. SPHERE will offer free access to all researchers and educators, for non-profit use. Users will be able to remotely access the resources, using their browsers and terminals. Users will obtain exclusive, on-demand access to resources they reserve and will be allowed to keep resources for a user-specified period of time, to access them as superusers and configure them as needed, to organize these resources into mini-networks, with configurable network substrate, and to experiment with malicious software as needed, in a safe, contained setting.

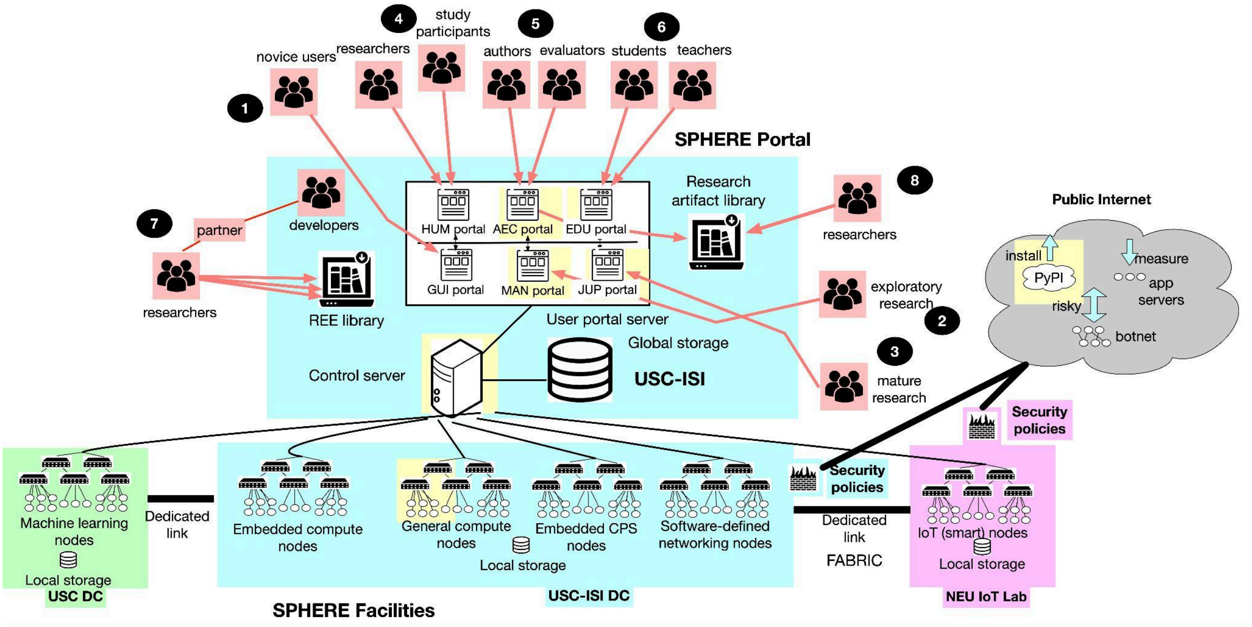

Figure 1: The SPHERE research infrastructure will offer access to an unprecedented variety of user-configurable hardware, software, and network resources. It will offer six user portals geared toward different populations of users. And, it will support reproducible research via infrastructure services as well as community engagement activities.

SPHERE facilities

At the lowest level in the architecture figure we show different classes of hardware nodes that support different research endeavors in cybersecurity and privacy, as summarized in Table 1.

Hardware	Description	Research Supported
General-compute	Around 200 server-class nodes, with Intel TDX, ARM CCA/TrustZone, and AMD SEV processor architectures General-compute resources can be requested as bare metal nodes or as virtual machines, with configurable resources	Research on application, system and network security, research that requires measurement of security and privacy phenomena in the Internet, research that uses human user studies, research that requires large-scale experimentation, and research that leverages trustworthy computing.
Embedded-compute	Around 400 nodes with embedded CPUs and GPUs, such as Intel Atom, Intel Xeon D, ARM Cortex-A57, and NVIDIA Jetson NX Volta	Research on edge computing security, blockchain security, private computing, trustworthy edge computing, and federated learning
Machine-learning	Around 10 servers with GPUs	Cybersecurity that includes machine-learning in the loop

Cyber-physical	Several complete architectures to emulate industrial control systems, such as a water-treatment plant, a power plant and an oil and gas pipeline	Security of critical infrastructure
Programmable	About 40 NetFPGA-equipped nodes	Research solutions that need dynamic (programmable) network security or solutions that investigate and improve SDN security
IoT	Around 500 IoT nodes (a variety of smart home, smart speaker, camera, doorbell, TV, appliance, medical, office, wearable, and miscellaneous devices)	IoT security and user privacy

Table 1: SPHERE will provide multiple facilities equipped with various classes of hardware nodes, supporting diverse initiatives within the cybersecurity and privacy research community.

Merge software

SPHERE facilities are powered by USC-ISI's Merge software for research infrastructure management, depicted in Figure 2.

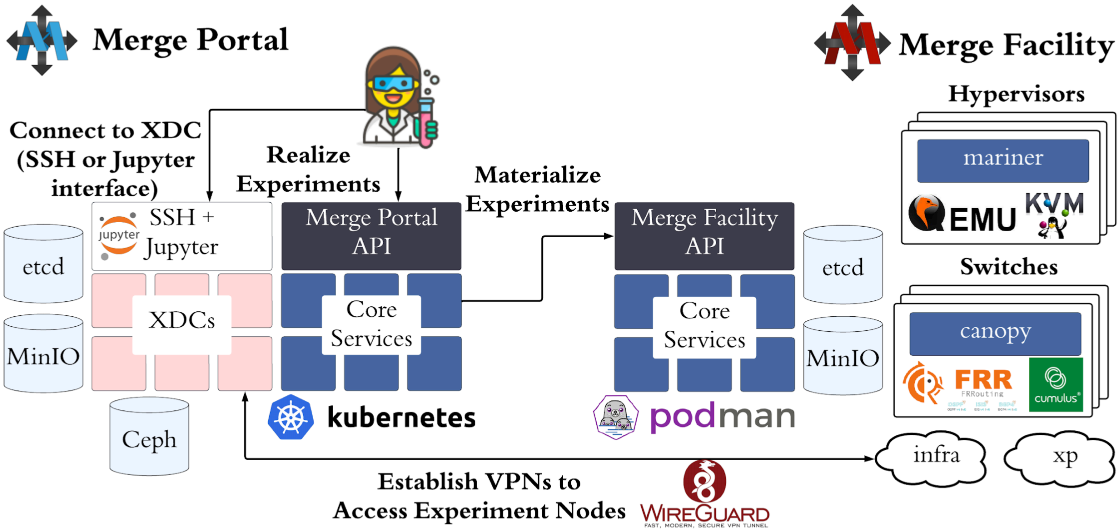

Figure 2: Merge research infrastructure software is designed to operate mid-scale testbeds with hundreds of compute nodes and tens of switches.

USC-ISI maintains reference implementations of the Merge Portal and Merge Facility code bases (https://gitlab.com/mergetb). The reference implementations use microservice architectures to flexibly integrate homegrown and third-party services to implement the Merge APIs. Both implementations target resilient operation at scale and strict adherence to user-defined performance requirements. The implementations were designed to operate mid-scale testbeds with hundreds of compute nodes and tens of switches, though they can be used to operate smaller scale testbeds if desired.

The ***Merge Portal*** implementation runs on Kubernetes, and provides two types of user services.

- ***Core services*** implement the Merge Portal API, and provide a range of services in support of compiling, realizing, and materializing experiments.
- ***XDCs*** (Experiment Development Container) are spawned as pods on the Kubernetes platform. They run lightweight Linux-based containers and provide a gateway to the testbed for users. XDCs connect to backend Merge facilities through on-demand VPN tunnels. Users connect to XDCs either through SSH or Jupyter (web-based HTTP).

The ***Merge Facility*** implementation runs on a set of different testbed resources.

- ***Infrastructure services*** host the Facility core services, testbed nodes providing virtual or bare metal device access to users, and switches that create virtual network segments to isolate user experiment traffic.
- ***Core services*** run as containers on top of podman, RedHat's container platform. These services implement the Merge Facility API, which the portal uses to send experiment requests to the backend facilities.
- ***Hypervisor services*** use Qemu/KVM to instantiate virtual machines for user experiments.
- ***Switches*** run on Cumulus Linux, a white-box switching platform provided by NVIDIA. The Merge canopy service runs on these switches to create isolated network segments through the VXLAN protocol.
- ***Dedicated infrastructure (infra)*** *and* ***experiment (xp)*** networks isolate testbed control traffic from experiment traffic to prevent interference between the two.

Merge supports multiple facilities, which may be managed by different teams and contain different hardware and software. Any compute/network infrastructure implementing the Merge Facility API can be commissioned as a Merge testbed facility.

SPHERE enclaves will be connected via dedicated Layer 2 links, including a FABRIC connection [8] between the IoT enclave at the Northeastern University and other enclaves at the USC-ISI and USC colocation facilities. Together, the Merge software and these connections will enable stitching of nodes from different enclaves into a single topology in an experiment.

SPHERE portal and security policies

The SPHERE portal, shown in the middle level of Figure 1, is hosted on distributed server-class nodes for resiliency, and enables all user access to SPHERE and enforce access policies.

At the same level, the right side of the figure illustrates different security policies that will be supported by SPHERE. Most experiments will be granted HTTP, HTTPS, and SSH access to the Internet, which is necessary for software installations and code downloads (e.g., from Github). Some experiments may need additional access, e.g., to facilitate Internet-wide measurements or risky interactions with malicious actors. These experiments will be supported through more open Internet access and additional, automated monitoring. Finally, some experiments may be so risky that they must be executed in full containment.

User portals

There are six user portals that cater to different user populations (shown at the upper level of Figure 1). Three basic portals are the manual, Jupyter, and graphical portals. The ***manual portal*** **(MAN)** enables direct access to experimental nodes via SSH, which facilitates exploratory research by expert users. When experimental workflows mature they can be scripted via the ***Jupyter portal*** **(JUP)** to allow for repeatable and reproducible experiments. Novice users will be offered access via the ***graphical user interface*** **(GUI)** ***portal***, allowing them to draw and annotate experiment topologies and workflows. Users will be able to switch between these portals as needed, keeping the experiment state.

SPHERE will also develop three specialized portals, to offer additional support to specific user populations. The ***education portal*** **(EDU)** enables teachers to create accounts for their students, to manage these accounts, to upload materials for class use, and to assign work to students to be completed on SPHERE. Such work usually comes in the form of homework assignments that require students to create attack and defense scenarios in mini-networks on SPHERE. This facilitates active learning, promotes student engagement and also teaches practical skills, which students will need in their future careers. Students also use the EDU portal to access materials for their class. The ***human study*** **(HUM)** ***portal*** helps researchers that run human user studies to deploy their innovations on SPHERE, and create pathways for study participants to interact with these innovations and leave feedback for researchers. The ***artifact evaluation committee (AEC) portal*** helps artifact authors and reviewers share and evaluate artifacts for a given paper on the same common infrastructure. Artifacts can then be archived for reuse by others.

Supporting reproducible experimentation

In addition to the physical architecture, SPHERE will offer a set of datasets and tools to facilitate representative, reproducible experimentation. First, it will enable and motivate users to package and archive their research artifacts into artifact libraries and make them available to others on the same platform. Second, it will actively work with artifact evaluation committees to support evaluation efforts on SPHERE and archive those artifacts that receive reproducibility badges. Third, SPHERE will crowdsource building of representative experimentation environments (REEs), which can serve as standards for evaluation in a given field of cybersecurity and privacy. SPHERE team will issue an annual call for mature research artifacts to be ported to SPHERE as REEs (please check our project page at https://sphere-project.net for this call). Artifact authors will receive summer funding to work as virtual interns and port their artifacts to SPHERE. Fourth, SPHERE will offer built-in support for artifact packaging and sharing, including support for experimental workflows, and recording of user actions during exploratory research (manual access to nodes), which can be used to transform exploratory experiments into mature ones by scripting the user’s manual actions.

3. Current Construction Status and Outreach

SPHERE construction has been ongoing for a year, and we are happy to report strong progress on all planned activities. We have procured and installed about one-third of the general-compute enclave and one-fifth of the IoT enclave. We have started design and purchasing of CPS, embedded compute and GPU enclaves. SPHERE portal and accompanying control, networking and storage infrastructure have been set up and SPHERE has officially opened to beta users in July 2024. Deterlab research and education users have also been migrated on to SPHERE and Deterlab has been decommissioned. The following table shows planned availability dates for different SPHERE enclaves.

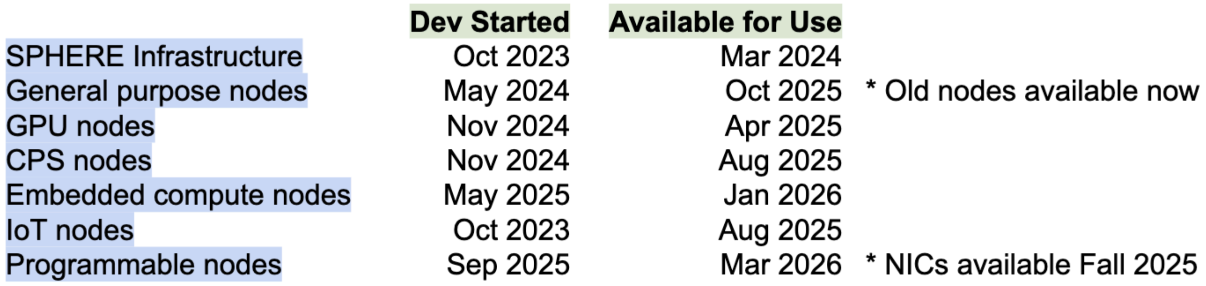

	Dev Started	Available for Use	
SPHERE Infrastructure	Oct 2023	Mar 2024	
General purpose nodes	May 2024	Oct 2025	* Old nodes available now
GPU nodes	Nov 2024	Apr 2025	
CPS nodes	Nov 2024	Aug 2025	
Embedded compute nodes	May 2025	Jan 2026	
IoT nodes	Oct 2023	Aug 2025	
Programmable nodes	Sep 2025	Mar 2026	* NICs available Fall 2025

Table 2: SPHERE enclaves will be developed and available for use according to a staggered basis over the first three years of the four-year project.

We have developed four out of six of the planned portals - MAN, JUP, EDU, and AEC. SPHERE is currently in use by more than 100 researchers and more than 600 students (10-12 classes) per semester. We welcome new beta users! You can join us at https://sphere-testbed.net.

As part of our community building and outreach efforts (spearheaded by Outreach Director David Balenson) we have engaged in extensive outreach at top cybersecurity conferences, symposia, and workshops, presenting posters and tutorials and leading birds-of-feather sessions. In 2024 we participated in the Internet Society Network and Distributed System Security (NDSS) Symposium, IEEE Symposium on Security and Privacy (S&P), IEEE European Symposium on Security and Privacy (EuroS&P), Cyber Security Experimentation and Test (CSET) Workshop, USENIX Security Symposium, ACM Conference on Computer and Communications Security (CCS), and Annual Computer Security Applications Conference (ACSAC). We also participated in two events that promote participation of underrepresented populations in computing, the CMD-IT/ACM Richard Tapia and SACNAS National Diversity in STEM (NDiSTEM) conferences.

We further participated in professional meetings that gather researchers in cybersecurity and privacy and in cyberinfrastructure: the NSF Cyber Innovation for Cyberinfrastructure (CICI) and Secure and Trustworthy CYberspace (SaTC PI) meetings, the NSF Research Infrastructure Workshop (RIW), Mid-scale Experimental Research Infrastructure Forum (MERIF), the NSF

Cybersecurity Summit, the FABRIC KNIT 8 workshop and Chameleon Community Workshop on Practical Reproducibility in HPC.

In addition to in-person meetings we engage in e-mail-based outreach to researchers, looking to understand their experimental needs. We hope to reach every potential user and offer them a chance to provide feedback on their needs and our current plans. Researchers and educators can also provide such feedback via a survey form that is linked to our project's page.

We have also worked with NDSS and Conference on emerging Networking EXperiments and Technologies (CoNext) artifact evaluation committees, and have lined up collaborations with several more for the second year of the project. We hosted eight paid summer interns to help us build SPHERE. These interns were recruited from institutions that serve a large number of minority and first-generation college students. Interns worked on a variety of tasks, including software and hardware installation and testing, front end and back end development, documentation, and automation.

4. Common vs. Private Research Infrastructure

Many researchers today experiment using private infrastructure, such as personal devices or devices in their research group's lab or their university. Even though some research solutions can be accurately evaluated in this setting, we argue that there are multiple reasons why evaluation on a common research infrastructure, such as SPHERE, brings substantial added benefits to the researcher and to the entire research community.

First, SPHERE will offer the scale and diversity of hardware, including most modern devices, that are beyond reach of many labs and university datacenters, along with a dedicated, responsive staff to provide user support.

Second, SPHERE will offer reproducibility support and processes to promote wider artifact sharing and reuse. Thus users that release their artifacts on SPHERE are likely to see these artifacts reused by others, increasing visibility and impact of their research.

Third, SPHERE will offer access to datasets, experimental tools, representative experimental scenarios and research artifacts shared by other users. These products will allow for easy experiment setup, where the researcher augments an existing, complex scenario instead of building the entire experiment from scratch. Artifact sharing further allows for head to head comparison between new and existing research products in same evaluation scenarios, which is necessary for research publications.

Finally, artifact sharing allows researchers to extend and enhance work of their peers, propelling the overall research community towards more sophisticated, more realistic and efficient solutions, which can more quickly transition to practice.

5. Conclusion

The SPHERE project is building a common, shared, community experimentation infrastructure for cybersecurity and privacy researchers and educators. We are excited to participate in its development, and we are working hard to learn about community research and education needs and incorporate these findings into our project plans. Please help us by collaborating with us —provide feedback, contribute research artifacts, become a beta user, or work as a SPHERE intern. This is your infrastructure—help us build it so we can all jointly benefit from it for many years to come!!

Acknowledgement & Disclaimer

SPHERE is based upon work supported by the National Science Foundation under Award #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.

References

[1] NPR. A 'Worst Nightmare' Cyberattack: The Untold Story Of The SolarWinds Hack. https://www.npr.org/2021/04/16/985439655/a-worst-nightmarecyberattack-The-untold-story-of-the-solarwinds-hack

[2] TechTarget.com. Colonial Pipeline hack explained: Everything you need to know. https://www.techtarget.com/whatis/feature/Colonial-Pipeline-hack-explained-Everything-you-need-to-know

[3] Statista. Annual number of ransomware attacks worldwide from 2016 to first half 2022. https://www.statista.com/statistics/494947/ransomware-attacks-per-yearworldwide/

[4] Government Technology. Hacktivism and DDOS Attacks Rise Dramatically in 2022. https://www.govtech.com/blogs/lohrmann-on-cybersecurity/hacktivismand-ddos-attacks-rise-dramatically-in-2022

[5] Sumeet Wadhwani, Spiceworks. Data Breaches Soared by 70% In Q3 2022 in an Otherwise Dull Year. https://www.spiceworks.com/it-security/datasecurity/news/data-breach-report/

[6] Balenson, D. et al. Cybersecurity artifacts workshop – report. https://bit.ly/CyberArtifactsWkshp2022

[7] Mirkovic, J., Balenson, D., Ravi, S., Garcia, L. & Benzel, T. Cybersecurity Experimentation Workshop – 2022 – Report. https://bit.ly/CyberExperWkshp2022

[8] FABRIC (website). https://portal.fabric-testbed.net/about/about-fabric

AUTHORS:

Jelena Mirkovic is a Principal Scientist and Project Leader at USC Information Sciences Institute and a Research Associate Professor at USC. She received her MS and PhD in Computer from UCLA and a BS in Computer Science and Engineering from the University of Belgrade, Serbia. Her research spans network-based attacks, human-centered attacks, cybersecurity experimentation, and Internet measurement. Her current research is focused on network security (botnets, denial-of-service attacks, and IP spoofing) as well as infrastructure, tools, and methodologies for conducting cybersecurity experimentation. She is the PI for the SPHERE project.

Brian Kocoloski is a Research Computer Scientist and Lead Scientist at USC Information Sciences Institute. He received his PhD in Computer Science from the University of Pittsburgh and a BS in Computer Science from the University of Dayton. His research focus is on developing secure, lightweight, and performant system software for networks, with applicability to a range of environments including data centers, clouds, and orbital systems in outer space. He leads the MergeTB project, a network testbed platform that supports a variety of advanced networking testbeds including the SPHERE testbed. He is a Co-PI and Technical Lead for the SPHERE project.

David Balenson is a Senior Supervising Computer Scientist and Interim Director of the Networking and Cybersecurity Division at USC Information Sciences Institute. He received his MS and BS in Computer Science from the University of Maryland. He has broad-based experience and background in critical infrastructure security and resilience, computer and network security, applied cryptography, and R&D project management. His current research interests include cybersecurity and privacy for critical infrastructure and cyber-physical systems including automotive and autonomous vehicles, experimentation and test, technology transition, and multidisciplinary research. He is the Community Outreach DIrector for the SPHERE project.