SPHERE: A Global Testbed for Reproducible AI, Cybersecurity, and Privacy Research

Jelena Mirkovic, David Balenson, and Erik Kline (USC-ISI), David Choffnes and Daniel Dubois (Northeastern University), Geoff Lawler, Joe Barnes, Yuri Pradkin, Christopher Tran, Srivatsan Ravi, Terry Benzel, and Alba Regalado (USC-ISI), Luis Garcia (U. Utah), and Ganesh Chennimalai Sankaran (RENCI)

Societal Need

- **Advancing research in cybersecurity and privacy** is of critical global importance for safeguarding people, infrastructure, and data worldwide
- As societies grow increasingly interconnected and reliant on digital systems, **robust and reproducible research** is essential to counter evolving threats and strengthen the security, privacy, and resilience of our shared global community

Research Need

- The global cybersecurity and privacy research community needs a **common, comprehensive, and representative research infrastructure** that meets the needs of all its members and enables reproducible science
- Such an infrastructure must support **realistic experimentation,** foster **widespread collaboration,** and accelerate the **development of solutions** that enhance cybersecurity and privacy worldwide

SPHERE Architecture and Capabilities

- **Diverse hardware to support diverse research needs (nearly 90% of today's publications):**
  - General and embedded compute nodes with trusted hardware, PLCs and IoT devices, programmable switches and NICs, and GPU-equipped nodes
- **Six user portals supporting:**
  - Exploratory research (MAN)
  - Novice users (GUI)
  - Mature research (JUP)
  - Education (EDU)
  - Human user studies (HUM)
  - Artifact evaluation (AEC)
- **Libraries of artifacts**
  - Realistic experimentation environments (REEs) and other artifacts
  - Easy reuse on SPHERE
- **Reproducibility support**
  - User action logging to alleviate cognitive load
  - Help package artifacts on SPHERE (including workflows)
  - Automatically verify completeness of an artifact
  - Support stability, consistency of results, and portability

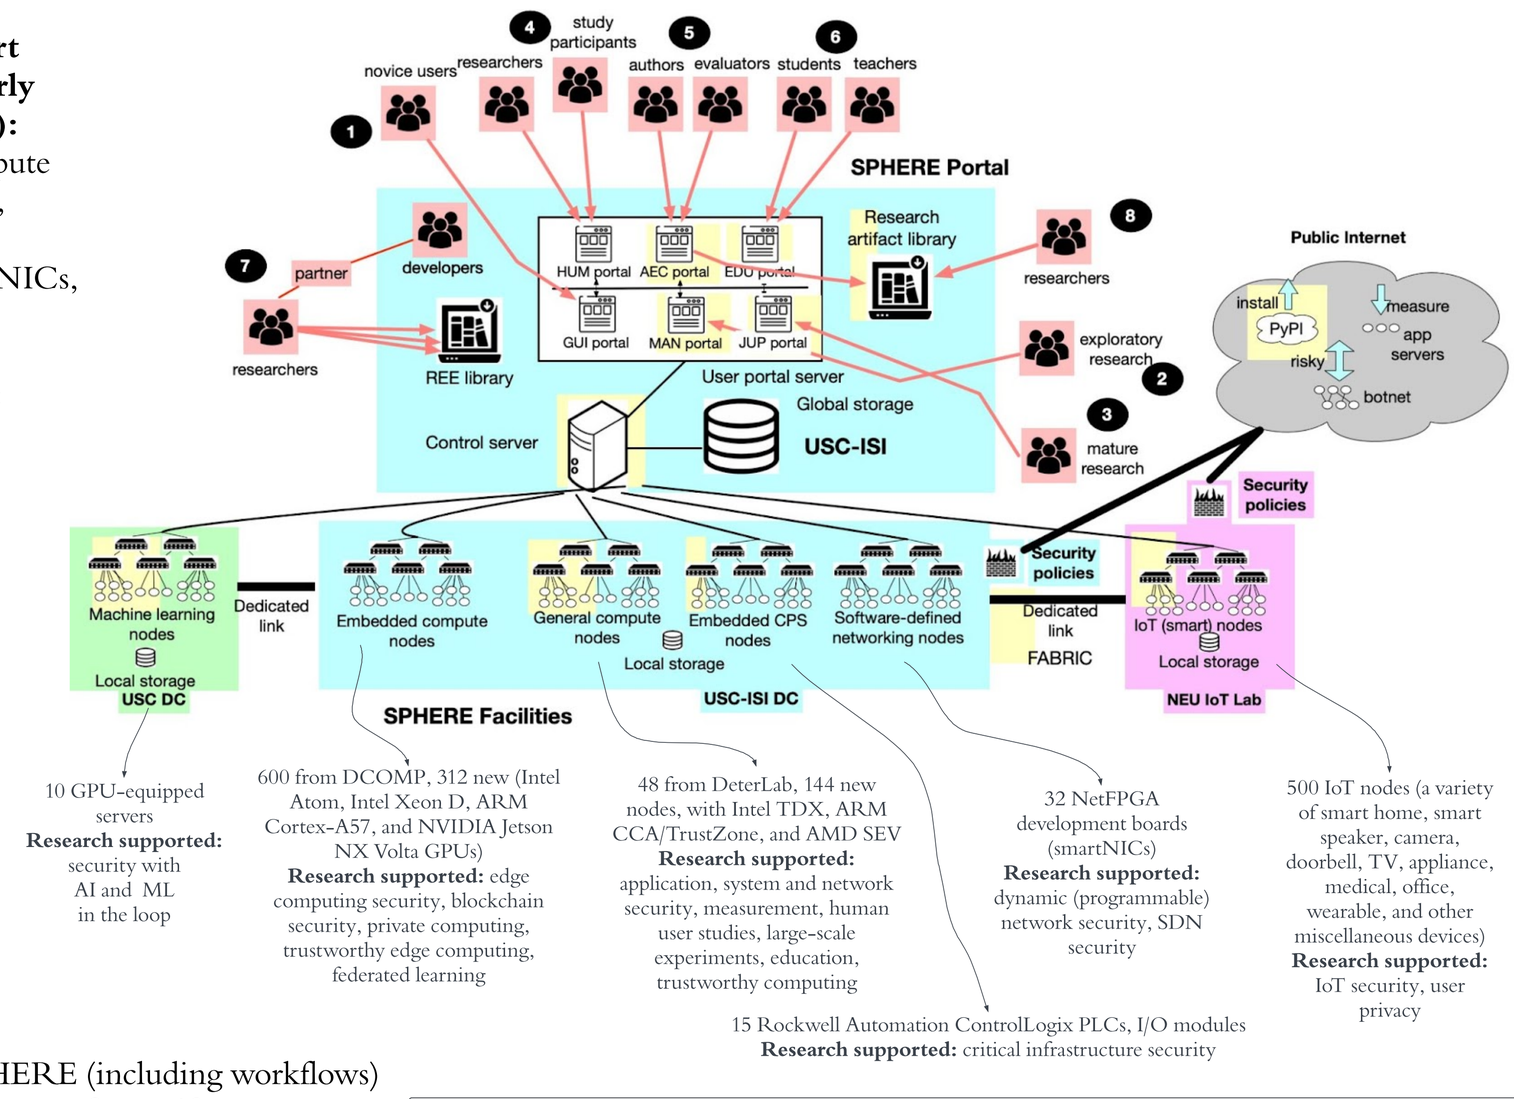

- **Flexible security policies:**
  - Full isolation
  - Measurement research
  - Software download
  - Risky experiments with malware
- **Sample use cases:**
  - Studying ICS security in a realistic environment
  - Studying IoT behavior and privacy implications
  - Studying AI-enhanced network attack detection and mitigation
  - Evaluation at different levels of fidelity
- Current/planned REEs:
  - Tier 1 ML security paper artifacts (UFL)
  - Tier 2 security paper artifacts (UFL)
  - CensorLab (UMass)
  - Reconstructed attack dataset (NU)
  - Heavy-duty vehicle CAN bus (CSU)
  - AI research support tools

Citation: Jelena Mirkovic, David Balenson, Brian Kocoloski. Enabling Reproducibility through the SPHERE Research Infrastructure. USENIX ;login: Online, USENIX Association. December 16, 2024. https://www.usenix.org/publications/loginonline/enabling-reproducibility-through-sphere-researchinfrastructure

SPHERE + ACCESS

SPHERE complements NSF ACCESS by supporting AI research workflows that require secure, high-fidelity, and risk-tolerant experimentation environments

- **Complementary infrastructure:** ACCESS provides large-scale shared compute and data resources; SPHERE provides isolated, instrumented environments for security-critical, adversarial, and privacy-sensitive AI experiments that are difficult to run on shared HPC systems
- **AI-in-the-loop experimentation:** SPHERE supports evaluation of AI and ML models operating within realistic networks, IoT deployments, CPS, and edge environments; SPHERE is exploring AI-driven experiment design, orchestration, and large-scale analysis
- **Secure and responsible AI research:** SPHERE enables controlled experiments involving malware, vulnerable systems, sensitive datasets, and privacy-preserving workflows, while maintaining strong isolation and reproducibility guarantees
- **Reproducibility and artifact support:** SPHERE provides libraries of reusable artifacts and REEs, automated checks for artifact completeness, and logging to support repeatable and transparent security, privacy, and AI research
- **Exploring integration pathways:** The project is exploring integration with national infrastructure programs such as NSF ACCESS and NAIRR, including pilot collaborations with researchers already using ACCESS resources

Current Status

- Well into third of four funded years
- Developing general-purpose, ML, IoT, CPS, and embedded enclaves
- Approx. one-third of general-purpose nodes and 100+ IoT devices available to beta users
- Designing prog. enclave
- Running control infrastructure and MAN, JUP, and EDU portals
- Piloting AEC portal, used for part of NDSS and S&P; working on artifact libraries
- Hosted 25 summer interns and several virtual interns
- Serving approx. 150 resarch beta users and approx. 1,000 educ. beta users

	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

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.