UK - US SEMICONDUCTOR SECURITY SYMPOSIUM

OCTOBER 2025

SUMMARY REPORT ON KEY RESEARCH CHALLENGES

Supported by the UK Department for Science, Innovation and Technology (DSIT)

UF UNIVERSITY of FLORIDA

RISE RESEARCH INSTITUTE FOR SECURE HARDWARE & EMBEDDED SYSTEMS

QUEEN'S UNIVERSITY BELFAST

INTRODUCTION

The 2025 UK-US Semiconductor Security Symposium convened experts from industry, government and academia from both the UK and the US at the Royal Society in London on the 8-9 October 2025. This briefing summarises the in-depth discussions held during the symposium to identify priority research challenges, laying the groundwork for collaborative research opportunities between US and UK research organisations. The four thematic areas discussed during the symposium were:

- Secure-by-Design Hardware Foundations
- AI-Enabled Security Design and Verification
- Supply Chain Integrity and Hardware Vulnerability Intelligence
- Open-Source Security IP and Ecosystem Cohesion

The 2025 symposium was the second in a series of collaborative UK-US semiconductor security events. The first workshop on *Security in the Era of Global Semiconductor Initiatives*, took place in November 2023 in Washington DC, US. That workshop culminated in a report published by the UK Research Institute in Secure Hardware and Embedded Systems (RISE) in July 2024 outlining a strategic overview of challenges and emergent opportunities in semiconductor security. The full report can be found on the RISE website: https://www.ukrise.org/

INSIGHTS & KEY RESEARCH CHALLENGES

SECURE-BY-DESIGN HARDWARE FOUNDATIONS

A central theme of the discussion on secure-by-design hardware foundations was the ambiguity surrounding what secure-by-design means in practice. Hardware engineers already consider security during design, yet attack costs continue to rise. This suggests that while security is present in design processes, it may not be systematically defined, measured, or prioritised. Participants distinguished between 'security by design' (explicit protective mechanisms embedded in the architecture) and 'security by default' (systematic reduction of attack surface), noting that both are relevant but neither is consistently operationalised.

There was agreement that secure-by-design must extend across the full product lifecycle, including internal and external supply chains, deployment, update mechanisms, and recovery. Recoverability and integrity restoration were identified as under-emphasised aspects of hardware security.

Participants also emphasised that the traditional performance–power–area (PPA) model is too narrow to capture the true trade-offs. Security must also be balanced against resilience, safety, forensics capability, scalability across platforms, and time-to-design. Importantly, the question was raised whether the long-standing assumption that security must be 'added without hurting performance' is itself limiting innovation. There was interest in reconsidering foundational architectural abstractions, in particular, CPU-centric digital processing models, and exploring whether alternative architectures or heterogeneous approaches might yield inherently stronger security properties.

In relation to standards and metrics, security is significantly harder to measure than performance or power. Participants highlighted the need for comparable, technology-agnostic metrics that focus on robustness rather than counting features. Analogies were drawn with energy ratings and automotive safety assessments, suggesting that graded or benchmarked approaches could improve transparency and accountability. However, it was recognised that measuring the absence of vulnerabilities is inherently difficult and that standardisation processes alone may be too slow without regulatory or market nudges.

A recurring concern was incentives. Performance optimisation is systematically rewarded in industry, while security often remains secondary unless driven by regulation, procurement requirements or reputational risk. Effective change may therefore require targeted regulatory frameworks, evaluation schemes, and stronger customer demand signals.

Key Research Challenges

- **Define secure-by-design operationally**: Establish a clear, lifecycle-wide and actionable definition of secure-by-design hardware.
- **Develop technology-agnostic robustness metrics**: Create measurable security and robustness metrics that enable comparison across architectures and platforms.
- **Advance scalable security verification**: Deliver scalable verification tools and methodologies capable of assessing complex and heterogeneous hardware systems.
- **Secure emerging architectures**: Investigate and strengthen security for chiplet-based systems and non-traditional compute paradigms.

AI-ENABLED SECURITY DESIGN AND VERIFICATION

AI is increasingly viewed as a necessary enabler in the face of growing design complexity and a shortage of skilled hardware security engineers. It offers the potential to act as an artificial security workforce, accelerating verification, red-teaming, and test generation.

AI-based approaches are already being used to explore test spaces more exhaustively and improve coverage. However, their effectiveness is constrained by limited availability of high-quality, trustworthy hardware security datasets. The absence of curated databases of secure designs, vulnerabilities, policies, and verification metadata was identified as a critical bottleneck.

Trustworthiness was highlighted as a central concern. Participants outlined the risks associated with hallucinations, lack of repeatability and hidden vulnerabilities in AI-generated outputs. Verification must therefore be holistic, incorporating threat modelling not only of hardware designs but also of the AI-enabled tools themselves. There was recognition that AI can both strengthen security and enable more sophisticated attacks, including automated vulnerability discovery and malicious functionality insertion.

The discussion also emphasised the need for open benchmarks, open hardware data, and open verification flows to enable reproducibility and validation. Certification and regulatory frameworks may be required to ensure reliability, particularly as AI becomes more deeply embedded in RTL generation, architectural design and large SoC integration.

Key Research Challenges

- **Develop AI-enhanced security design and verification tools**: Advance AI techniques that strengthen design, testing, verification, and red-teaming across hardware development flows.
- **Build trusted security datasets**: Create curated, trustworthy databases of secure designs, vulnerabilities, policies and verification artefacts to enable effective AI training and benchmarking.
- **Advance AI-driven threat modelling and attack generation**: Develop generative-AI approaches for automated threat modelling, attack scenario generation and coverage assessment.
- **Secure AI-enabled and edge-AI systems at the hardware level**: Investigate the hardware security implications of AI accelerators, edge-AI deployments and AI-based SoC integration.
- **Design certification and regulatory frameworks for AI in hardware design**: Create evaluation mechanisms that ensure reliability and accountability of AI-assisted semiconductor design tools.

AI can both strengthen security and enable more sophisticated attacks, including automated vulnerability discovery and malicious functionality insertion.

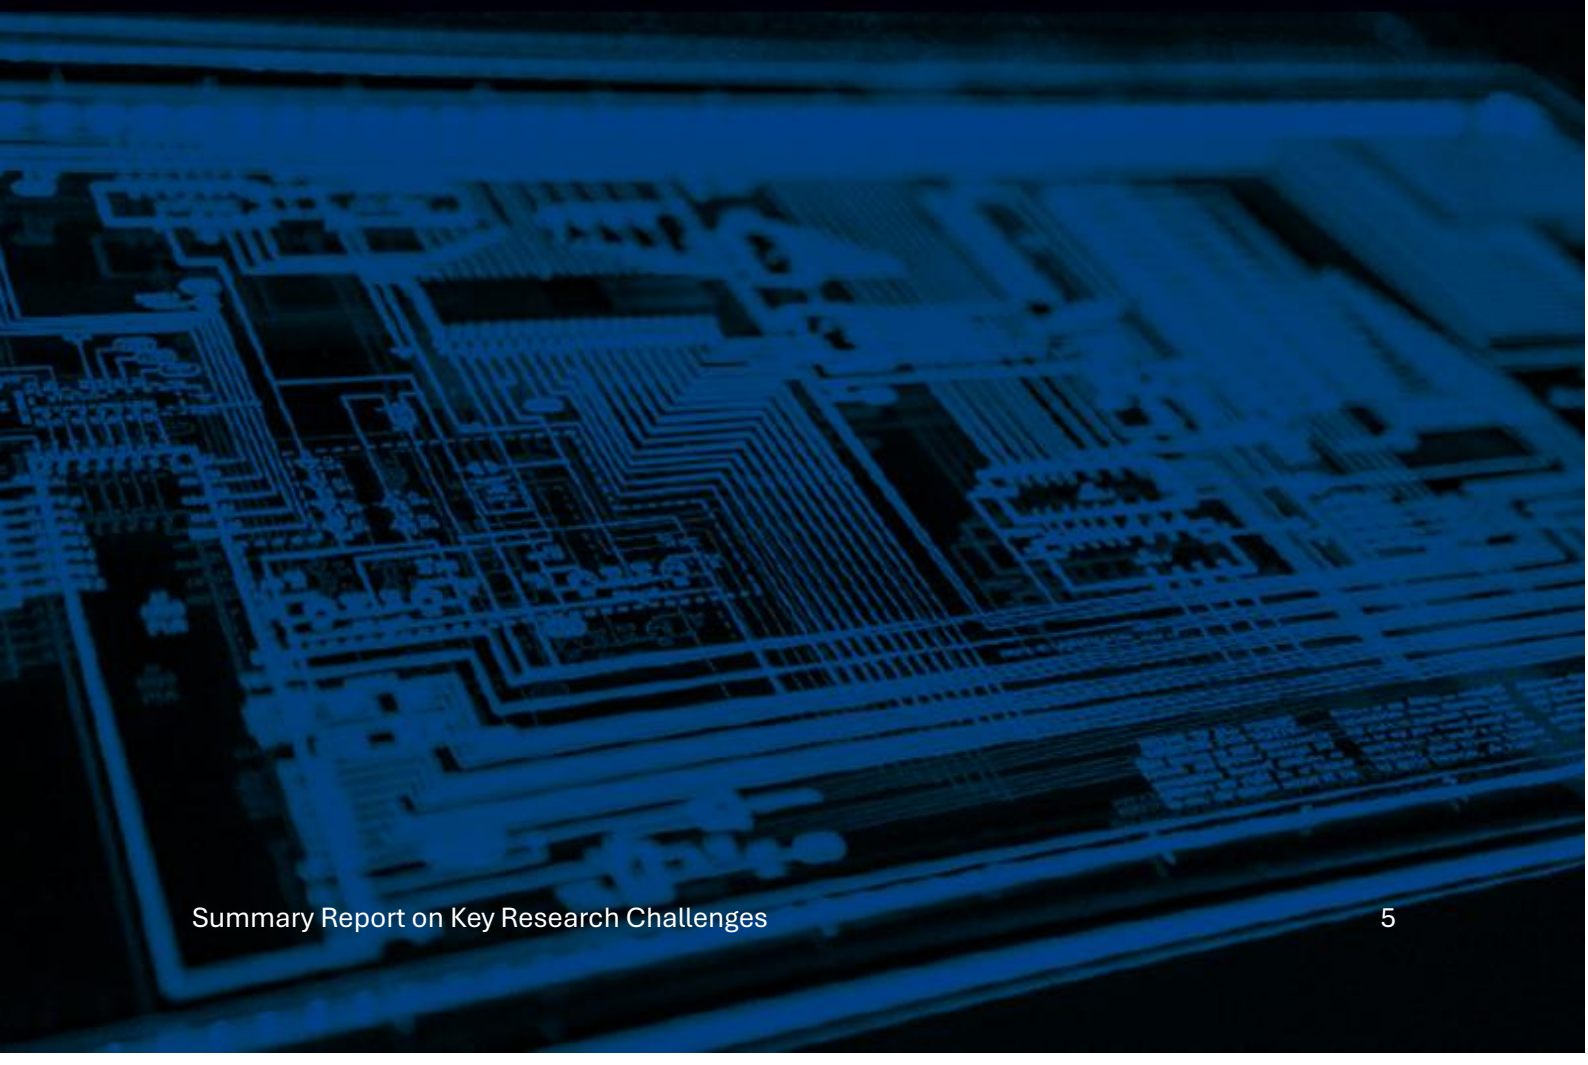

SUPPLY CHAIN INTEGRITY AND HARDWARE VULNERABILITY INTELLIGENCE

There was broad support for a hardware vulnerability database; however, caution was urged against duplicating existing efforts. Rather than creating an entirely new database, participants suggested extending established frameworks (e.g. CWE/CVE-style approaches) to incorporate hardware-specific artefacts, reproducibility data and supply chain context. A key requirement is improved classification, authentication levels and confidence scoring for reported vulnerabilities. Mapping vulnerabilities to explicit risk levels was seen as essential to drive behavioural and procurement change.

Participants noted that vulnerability intelligence is most valuable during component selection and system design, but considerably harder to operationalise for already deployed hardware. Verification and qualification of hardware threats require significant effort and resources, and voluntary disclosure models may limit completeness unless supported by incentives or regulatory drivers.

With respect to traceability and provenance, there was strong agreement on continued relevance. Hardware supply chains, particularly in the context of heterogeneous integration and chiplet-based designs, introduce new attack surfaces and complexity. Unlike software SBOM (Software Bill of Materials) approaches, hardware cannot rely solely on simple hashing mechanisms; pre-fabrication and post-fabrication assurance require different technical strategies. Standardised root-of-trust mechanisms, supply chain authentication schemes and potentially distributed ledger-style approaches were discussed as possible enablers.

Cost and scalability of supply chain audits emerged as a significant concern, particularly for smaller organisations. Participants highlighted the need to distinguish between company-level assurance and component-level verification, and to explore trust-grading or reputation-based mechanisms to support ecosystem-wide assurance.

Key Research Challenges

- **Integrate and extend hardware vulnerability frameworks**: Build upon existing databases to incorporate hardware-specific weaknesses, reproducibility artefacts, classification confidence levels, and explicit risk mappings.
- **Develop vulnerability severity metrics**: Establish metrics that remain robust over time and support risk-informed procurement and lifecycle management decisions.
- **Engineer lifecycle-wide traceability mechanisms**: Design authentication and provenance techniques spanning design, fabrication, integration, deployment and maintenance stages.
- **Secure heterogeneous and chiplet-based ecosystems**: Develop threat models and secure root-of-trust architectures for multi-party integration environments.
- **Protect AI and model supply chains**: Develop methods to detect and prevent poisoning, tampering and unauthorised modification within hardware and AI model supply chains.

OPEN-SOURCE SECURITY IP AND ECOSYSTEM COHESION

While open-source hardware can act as a shared resource toward common goals and enable transparency, education and research, its quality is uneven and cannot be assumed to confer security by default. Participants cautioned against the belief that 'many eyes' automatically improve security, noting that open artefacts are equally visible to attackers. At the same time, closed-source solutions are not inherently secure; quality assurance and governance are decisive factors.

A key tension identified was the limited openness of commercial EDA ecosystems and the persistence of siloed IP models, which can hinder shared security capability development. Open platforms were viewed as essential for education and exploratory research, particularly in academia and early-stage innovation. Examples of emerging open roots of trust and industry-adopted open cores were discussed as positive signals, though sustainability and governance remain critical concerns.

Standards and metrics were also discussed under this theme and participants expressed scepticism about poorly defined security metrics. There was recognition that verification plans often test articulated properties, while unarticulated or emergent properties remain harder to assess. The concept of 'hardware contracts' around emergent security properties was suggested as a possible future direction.

Overall, ecosystem cohesion involving sustainable governance models, shared verification practices and common metrics was viewed as necessary to build trust and avoid fragmentation.

Priority Research Challenges

- **Strengthen open-source hardware security foundations**: Develop sustainable governance, quality assurance, and maintenance models for open-source security IP.
- **Standardise security verification practices**: Develop common verification methodologies and artefact-sharing mechanisms aligned with industry workflows.
- **Bridge open and commercial ecosystems**: Design collaboration models that enable interoperability between open-source initiatives and proprietary EDA and IP environments.
- **Enable open platforms for education and research**: Ensure access to open hardware platforms and toolchains to support skills development and exploratory security research.

CONCLUSIONS

The 2025 UK-US Semiconductor Security Symposium reaffirmed that semiconductor security is now a strategic, systemic and transnational priority. Participants agreed that security must be embedded across the full lifecycle, supported by measurable robustness, trustworthy AI-enabled tools and coherent ecosystem governance.

A key insight was that the challenges are not solely technical. Progress in verification, metrics, AI-assisted design and supply chain assurance must be matched by stronger incentives, standards, procurement practices and regulatory alignment. Security must be treated as a first-order design parameter rather than secondary to performance and cost.

Discussions also highlighted the dual-use nature of emerging technologies: AI can strengthen verification but also enable more advanced attacks; open-source hardware promotes innovation but does not guarantee security; and supply chain transparency must be scalable and economically viable to be effective.

The symposium underscored the importance of sustained UK–US collaboration. Coordinated development of datasets, benchmarks, verification frameworks and certification approaches would accelerate progress and reduce fragmentation. The key research challenges identified in this report provide a structured agenda for long-term bilateral cooperation, aligning research and policy to shape a resilient and globally relevant model for secure semiconductor innovation.

RESEARCH INSTITUTE FOR
SECURE HARDWARE &
EMBEDDED SYSTEMS

CONTACT DETAILS
W: www.ukrise.org
E: info@ukrise.org

ukrise