NIST Internal Report
NIST IR 8615

Workshop on Rolling Next-Generation Secure Hardware into Standards

Sanjay (Jay) Rekhi
Nedim Goren
Yang Guo

This publication is available free of charge from:
https://doi.org/10.6028/NIST.IR.8615

NIST NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY
U.S. DEPARTMENT OF COMMERCE

NIST Internal Report
NIST IR 8615

Workshop on Rolling Next-Generation Secure Hardware into Standards

Sanjay (Jay) Rekhi
Nedim Goren
Yang Guo
Computer Security Division
Information Technology Laboratory

This publication is available free of charge from:
https://doi.org/10.6028/NIST.IR.8615

September 2026

U.S. Department of Commerce
Howard Lutnick, Secretary of Commerce

National Institute of Standards and Technology
Arvind Raman, NIST Director and Under Secretary of Commerce for Standards and Technology

Certain equipment, instruments, software, or materials, commercial or non-commercial, are identified in this paper in order to specify the experimental procedure adequately. Such identification does not imply recommendation or endorsement of any product or service by NIST, nor does it imply that the materials or equipment identified are necessarily the best available for the purpose.

There may be references in this publication to other publications currently under development by NIST in accordance with its assigned statutory responsibilities. The information in this publication, including concepts and methodologies, may be used by federal agencies even before the completion of such companion publications. Thus, until each publication is completed, current requirements, guidelines, and procedures, where they exist, remain operative. For planning and transition purposes, federal agencies may wish to closely follow the development of these new publications by NIST.

Organizations are encouraged to review all draft publications during public comment periods and provide feedback to NIST. Many NIST cybersecurity publications, other than the ones noted above, are available at https://csrc.nist.gov/publications.

NIST Technical Series Policies
Copyright, Use, and Licensing Statements
NIST Technical Series Publication Identifier Syntax

Publication History
Approved by the NIST Editorial Review Board on 2026-07-20

How to Cite this NIST Technical Series Publication
Rekhi S, Goren N, Guo Y (2026) Workshop on Rolling Next-Generation Secure Hardware into Standards. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Internal Report (IR) NIST IR 8615. https://doi.org/10.6028/NIST.IR.8615

Author ORCID iDs
Sanjay (Jay) Rekhi: 0009-0008-8711-4030
Ned Goren: 0009-0009-3578-2958
Yang Guo: 0000-0002-3245-3069

Contact Information
hwsec@nist.gov

National Institute of Standards and Technology
Attn: Computer Security Division, Information Technology Laboratory
100 Bureau Drive (Mail Stop 8930) Gaithersburg, MD 20899-8930

Additional Information
Additional information about this publication is available at https://csrc.nist.gov/pubs/ir/8615/final, including related content, potential updates, and document history.

All comments are subject to release under the Freedom of Information Act (FOIA).

Abstract

The Sustainable Hardware Security (SUSHI@NIST) Workshop on January 26, 2026, convened stakeholders from industry, academia, and government to identify priorities and develop actionable recommendations for advancing hardware security across the semiconductor ecosystem. This report summarizes the workshop's outcomes, including a call to develop unified, interoperable, and evidence-based trust models, vocabularies, and guidance on governing hardware security. Public-private partnerships, pilot programs, and aligning incentives through appropriate procurement requirements can reduce economic barriers and demonstrate secure, provenance-enabled semiconductor ecosystems.

Keywords

AI hardware security; embedded systems security; firmware security; hardware security; life cycle security; memory security; microelectronics security; scalable security assurance; semiconductor assurance; semiconductor security standards; side-channel analysis.

Reports on Computer Systems Technology

The Information Technology Laboratory (ITL) at the National Institute of Standards and Technology (NIST) promotes the U.S. economy and public welfare by providing technical leadership for the Nation's measurement and standards infrastructure. ITL develops tests, test methods, reference data, proof-of-concept implementations, and technical analyses to advance the development and productive use of information technology. ITL's responsibilities include the development of management, administrative, technical, and physical standards and guidelines for the cost-effective security and privacy of other than national security-related information in federal information systems.

Table of Contents

Acknowledgments

The authors would like to thank the workshop organizing committee, presenters, panel members, and participants who provided valuable input via questions, answers, and the breakout discussions.

1. Introduction

Semiconductors are a fundamental part of the modern global economy and a key enabler of national security, consumer technology, and advanced infrastructure. Chips are the digital backbone of nearly all essential systems, including artificial intelligence (AI), transportation, telecommunications, healthcare, and defense platforms. The rapid increase in semiconductor demand — driven in part by AI and data-intensive workloads — has heightened the strategic importance of secure, resilient, and trusted microelectronics across commercial and government sectors.

Recognizing its importance, the United States has launched a comprehensive national strategy to strengthen the semiconductor ecosystem and maintain technological leadership. Federal investments are fast-tracking domestic manufacturing, boosting research and development, and building a skilled workforce. U.S.-based companies still lead globally in revenue and innovation. Furthermore, manufacturing capacity is being reshored and expanded domestically. These efforts are supported by trade requirements, incentives, and international partnerships that are designed to enhance supply chain resilience and reduce reliance on geographically concentrated production.

Within this broader approach, cybersecurity has become a key pillar of semiconductor security. Modern threats go beyond traditional software vulnerabilities to include risks embedded in hardware and the supply chain, such as intellectual property theft, counterfeit components, malicious logic inserted during design or manufacturing, and disruptions to fabrication facilities. Since semiconductors are essential to critical infrastructure and defense systems, a compromise at any stage of the life cycle can have cascading effects on national security, economic stability, and operational continuity.

The National Institute of Standards and Technology (NIST) plays a key role in advancing semiconductor cybersecurity through guidelines development, metrology, and collaborative research. Securing semiconductors requires a comprehensive view across three interconnected areas: (1) design security; (2) manufacturing security, including fabrication facilities and outsourced assembly and test (OSAT) environments; and (3) supply chain security to ensure the origin and traceability of components. Specific initiatives (e.g., semiconductor-specific cybersecurity profiles, hardware assurance programs, traceability research, CHIPS Program safeguards) help create a complete framework for trusted microelectronics.

Beyond NIST, various coordinated efforts across government and industry are strengthening semiconductor security. The U.S. Department of Defense supports initiatives like the Microelectronics Commons to accelerate secure prototyping and move technologies from research to production. The Defense Advanced Research Projects Agency (DARPA) is advancing next-generation hardware security and manufacturing methods, including heterogeneous integration approaches that boost performance while making it harder for adversaries to access. Industry groups like the SEMI Manufacturing Cybersecurity Consortium are working to align global standards and improve supply chain practices. Overall, these efforts show a common understanding that semiconductor security is vital for safeguarding intellectual

property and manufacturing capabilities as well as maintaining trust, resilience, and strategic advantage in an increasingly digital and interconnected world.

1.1. SUSHI@NIST Workshop Details

As part of NIST's broader efforts to enhance semiconductor security, the Sustainable Hardware Security (SUSHI@NIST) workshop was held virtually on January 26, 2026. The workshop brought together stakeholders from government, industry, and academia to address security challenges across the semiconductor development life cycle. The workshop focused on threat analysis, mitigation strategies, and aligning emerging standards with practical operational needs with goals to promote cross-sector collaboration, help identify gaps in current practices, build consensus on best practices, and shape future standards and guidance to support trusted hardware ecosystems.

The SUSHI platform was envisioned and managed by:

- Gang Qu, University of Maryland
- Jeyavijayan Rajendran, Texas A&M
- Ahmad-Reza Sadeghi, TU Darmstadt

2. Workshop Details

The SUSHI@NIST workshop was held virtually with a total of 116 participants. Of those, 65% represented industry, 20% represented government, and 15% represented academia. The workshop agenda is included in Appendix A.

The agenda and the presentations that were allowed to be made publicly available can also be found at: https://www.nist.gov/news-events/events/2026/01/sushinist-rolling-next-generation-secure-hardware-standards.

2.1. Speaker Viewpoints

2.1.1. A Roadmap Toward Scalable Security Assurance

Jason Fung outlined the growing challenge of achieving scalable security assurance in modern hardware systems and emphasized that increasing system complexity — especially in semiconductors and systems-on-a-chip (SoCs) , has made traditional, manual approaches insufficient. Today's chips contain thousands of components, diverse IP blocks, and intricate hardware-software interactions, all of which introduce significant verification burdens. Security assurance is not simply the absence of vulnerabilities but the result of proactive investment in prevention, detection, and remediation across the product life cycle.

Fung argued that security assurance must scale across architectures, teams, and toolchains, despite fragmentation in design environments and expertise. There is no one-size-fits-all method; instead, assurance requires layered, complementary techniques that span pre-silicon verification, post-silicon testing, supply chain security, and long-term deployment monitoring.

Fung proposed a structured, scalable framework analogous to how compilers transformed software development. By formalizing principles (e.g., secure by design/default), processes (e.g., end-to-end life cycle integration), and guidance (e.g., standards, compliance, AI guardrails), security assurance can be increasingly automated and embedded into development workflows. However, much of today's expertise relies on institutional knowledge, creating barriers to scale and workforce development.

Fung also highlighted the need for industry-wide collaboration, including leveraging existing software security frameworks, developing reusable reference designs, and investing in tools, taxonomies, and training programs. Initiatives like hardware vulnerability databases, academic partnerships, and workforce training (e.g., CTFs, certification programs) demonstrate progress but require broader adoption. Ultimately, scalable assurance depends on aligning people, processes, and technology with the support of community-driven standards and sustained research and development (R&D).

2.1.1.1. Call to Action/Next Steps

- Define a shared roadmap and governance model

Establish a multi-year vision for scalable hardware security assurance, including clear stakeholder roles across industry, academia, and government. Align on mission, priorities, and measurable outcomes to guide coordinated progress.

- Standardize frameworks and best practices

  Formalize secure-by-design principles, extend proven software frameworks (e.g., threat modeling, attack taxonomies) to hardware, and develop common guidance to reduce fragmentation and improve consistency across organizations.

- Invest in automation, tools, and R&D

  Accelerate the development of scalable verification techniques (e.g., AI-assisted analysis, static analysis, fuzzing) and fund research to close gaps. Incentivize innovation through grants, challenges, and ecosystem partnerships.

- Build workforce capacity and knowledge sharing

  Bridge the gap between academia and industry by integrating practical security training, expanding collaborative programs, and codifying institutional knowledge into accessible resources to lower barriers for new practitioners.

2.1.2. Security in Memory and Storage

Tamara Schmitz highlighted that memory, once viewed as a commodity, is now a vital and increasingly targeted attack surface because of its central role in AI and modern computing. Memory (e.g., DRAM, HBM) temporarily stores data in use, while storage (e.g., SSDs, NAND) keeps data at rest, resulting in distinct but related security issues. Attackers are exploiting both through hardware-level techniques like row hammer, cold boot, direct memory access, and advanced side-channel and fault injection attacks. Worryingly, the cost and availability of these attacks are dropping rapidly, and various tools (e.g., low-cost laser fault injection setups, AI-driven fuzzing) make sophisticated exploits widely accessible.

Schmitz discussed how storage introduces additional risks due to firmware, including software vulnerabilities (e.g., buffer overflows) and supply chain threats. Across both memory and storage, attackers exploit debug interfaces, metadata leakage, and weaknesses in life cycle handling. While mitigations exist (e.g., encryption, hardware roots of trust, firmware signing, error correction), their implementations have been inconsistent and often driven by market demands rather than standardized baselines. Current standards are lagging behind evolving threats and are often ambiguous when applied to hardware, which can lead to inconsistent risk assessments.

A major concern is the industry's reactive stance: vulnerabilities are often addressed only after they have been exploited. Schmitz advocates for a proactive, life cycle-based security approach from design to end of life, along with stronger coordination among hardware and firmware teams. She also highlights the importance of third-party assurance, secure development life cycles, and supply chain transparency, including SBOMs and contractual security requirements. Ultimately, she urges clearer and more uniform standards, better threat modeling, and ongoing

adaptation to new attack methods. Security must be integrated early in the design process and maintained throughout the product life cycle with collaboration among vendors, partners, and customers.

2.1.2.1. Call to Action/Next Steps

- Establish baseline security standards for memory and storage, including essential features (e.g., attestation, rollback protection, consistent encryption practices). This clarifies expectations and guarantees a basic level of security, regardless of vendor or market.
- Shift from reactive to proactive security by integrating protections early in design and maintaining them throughout the product life cycle. Regularly review threat models, test against new attacks, and plan for long-term support and patching.
- Improve supply chain and third-party security by implementing secure development life cycle practices, requiring SBOMs, and including security requirements in contracts. Ensure consistency across all partners and components.
- Improve cross-domain coordination and incident response by integrating hardware and firmware teams, expanding threat intelligence capabilities, and enabling faster and more unified responses to complex, multi-layer attacks.

2.1.3. Hardware Attacks on AI Inference Chips

Marc Witteman from Keysight Technologies presented early research on hardware attacks that target AI inference chips, which differ from traditional cryptographic targets and introduce new security risks. He explained the difference between AI training (i.e., resource-intensive and centralized) and inference (i.e., lightweight deployment on edge devices), emphasizing that trained model weights are valuable intellectual property and that inference integrity is vital for safety-critical decisions (e.g., autonomous vehicles). Using a YOLO-based ("You Only Look Once") object detection model on an AI-enabled chip, Witteman showed how side-channel analysis (e.g., electromagnetic emanations) uncovers subtle processing differences, depending on input data. These signals help identify when and where inference computations occur, enabling targeted attacks. He also demonstrated a brief voltage fault injection to disrupt power and manipulate the chip's behavior.

Witteman's results showed that hardware glitches can meaningfully alter AI outputs. About 77% of experiments caused shifts in confidence scores, which could lead to missed detections (e.g., failing to recognize a pedestrian). More importantly, around 22% produced "hallucinations," where the system detected objects that do not exist, sometimes with high confidence. These effects could cause unsafe actions in real-world systems (e.g., autonomous vehicles or robotics). Witteman concluded that AI inference chips need strong hardware security protections to ensure both model confidentiality and decision accuracy. Although current results are somewhat random, future research aims to produce precise, attacker-

controlled outcomes. He also warned about supply chain risks, where malicious hardware modifications could embed attack capabilities directly into deployed systems.

2.1.3.1. Call to Action/Next Steps

- Advance hardware security for AI chips

  Develop and integrate hardware-level protections (e.g., fault detection, shielding, secure power design) that are specifically tailored for AI inference workloads to protect both model weights and decision integrity.

- Expand research on controlled attacks

  Move from probabilistic glitching to deterministic manipulation, enabling a deeper understanding of attacker capabilities and informing realistic threat models and defenses.

- Establish testing and validation standards

  Develop standardized methods for assessing AI hardware resilience to side-channel and fault injection attacks, similar to established standards in cryptographic hardware security.

- Address supply chain threats

  Investigate the risks of malicious hardware insertion (e.g., mod chips) and develop provenance, inspection, and verification mechanisms to ensure the trusted deployment of AI-enabled systems.

2.1.4. NIST Semiconductor Standards Activities

Mary Bedner outlined NIST's semiconductor standards activities within the CHIPS for America Program, emphasizing that her overview reflects key efforts led by the CHIPS R&D office rather than a complete inventory.

NIST's strategy centers on creating a smarter, faster, and more agile semiconductor standards ecosystem aligned with global principles. Key priorities include enhancing private-sector leadership, speeding up standards development, increasing awareness and ease of navigation for standards, and facilitating global market access. A significant milestone was the 2023 standards summit, which brought together over 600 stakeholders to pinpoint ecosystem-wide priorities, such as chiplets, supply chain security, advanced packaging, and faster standards innovation. This was followed by targeted workshops that focused on chiplet interfaces, digital twins, and data standards, highlighting ongoing challenges in data sharing and interoperability.

Internally, NIST coordinates through interagency working groups and the CHIPS metrology program, which promotes pre-standards R&D in areas like measurement science, traceability, and hardware security. Current projects include advanced packaging metrology, hardware vulnerability management, and secure data sanitization.

To drive progress, NIST launched a Broad Agency Announcement (BAA) to fund R&D, commercialization, and standards development in emerging technologies like AI, quantum computing, and biotechnology. Additionally, a new public-private partnership with ASTM International created the ASCET (Advancing Standardization for Critical and Emerging Technologies). The Center of Excellence enhances U.S. leadership in global standardization, focusing on workforce development, collaboration, and data-sharing infrastructure. Bednar emphasized that stakeholder engagement is essential for turning identified priorities into widely adopted standards.

2.1.4.1. Call to Action/Next Steps

- Advance identified priorities into standards

  Review published workshop and summit outputs and collaborate with standards development organizations to convert priority areas (e.g., chiplets, supply chain security, data standards) into formal standards or industry consortia outputs.

- Engage with CHIPS R&D funding opportunities

  Submit proposals to the BAA program, particularly in standards development and emerging technology domains, while addressing ROI considerations and aligning with national priorities.

- Participate in ecosystem coordination efforts

  Contribute to interagency reports, workshops, and public-private initiatives (e.g., ASCET Center) to reduce silos, improve standards visibility, and strengthen cross-sector collaboration.

- Invest in workforce and data infrastructure

  Support the development of a skilled standards workforce and engage in initiatives that enhance data sharing, interoperability, and traceability across the semiconductor ecosystem.

2.1.5. NSF/CISE Security, Privacy, and Trust in Cyberspace (SaTC 2.0)

Selcuk Uluagac provided an overview of National Science Foundation (NSF) activities. The NSF's SaTC 2.0 program is a flagship initiative that supports interdisciplinary, foundational research in cybersecurity, privacy, and trust worldwide. Since it started in 2011, the program has invested over $1 billion in nearly 4,000 awards, promoting advancements in resilience, privacy-preserving technologies, and societal impacts of cyber systems. It highlights collaboration across fields like computer science, social sciences, economics, and policy, working closely with other federal agencies and international partners.

SaTC 2.0 reflects the changing cyber landscape and incorporates emerging technologies like AI, quantum computing, and large-scale distributed systems. The program emphasizes "trustworthy by design" approaches and encourages researchers to rethink security from first

principles instead of retrofitting existing systems. It also highlights AI's dual role as both a tool to improve security and a possible attack vector. Key research areas include hardware security, adversarial machine learning, information integrity, human-centered security, and privacy-preserving data systems. Recent updates to the solicitation simplify proposal categories and introduce new funding structures, such as "Research" and "Seed" designations. The program promotes proposals that are aligned with national strategies and focuses on broader impacts, including workforce development and increasing participation in computing. Engagement with the research community through regular workshops, PI meetings, and outreach activities remains essential.

Overall, SaTC 2.0 positions NSF as a foundational funding body at low technology readiness levels (TRLs), enabling early-stage, high-impact research that feeds into applied efforts across government and industry.

2.1.5.1. Call to Action/Next Steps

- Align proposals with SaTC 2.0 priorities

  Focus on interdisciplinary, "trustworthy by design" research that addresses emerging challenges, such as AI security, hardware trust, and information integrity. Ensure alignment with national cybersecurity and AI strategies to improve competitiveness.

- Leverage collaboration opportunities

  Engage with academic, industry, and international partners. Consider cross-disciplinary teams and utilize NSF-supported research infrastructure (e.g., cloud labs, AI resources) to strengthen proposals.

- Prepare for updated proposal requirements
- Understand new designations (i.e., Research, Seed), funding limits, and requirements, such as Broadening Participation in Computing (BPC) plans for larger projects. Monitor target dates (January and September) while noting that rolling submissions are allowed.
- Engage with NSF and its community early

  Participate in webinars, PI meetings, and outreach events. Reach out to program directors for guidance, explore reviewer roles, and stay informed on Dear Colleague Letters for emerging "Seed" topics.

2.1.6. Standards-Based Assurance for Trusted Electronics

Jim Will emphasized the vital need for standards-based assurance to establish trust in global electronics supply chains, especially for defense and critical infrastructure. He pointed out that while hardware assurance challenges have existed for a long time, they are growing more serious due to complex, worldwide supply chains and heavy dependence on offshore manufacturing, particularly in Asia. Current visibility into supply chains is limited, often only

extending to top-tier suppliers, which leaves hidden risks (e.g., foreign dependencies, potential adversarial exposure).

Through case studies (e.g., military-use tablets and naval systems), he illustrated how widely used commercial-off-the-shelf (COTS) components — although cost-effective and scalable — often lack sufficient assurance. These systems can unknowingly incorporate parts from untrusted or restricted sources. In contrast, custom or high-assurance microelectronics (e.g., secure FPGAs, trusted foundry flows) offer stronger security but are less scalable and more costly. Will argued that standards and domestic or allied manufacturing capabilities must develop together. Fragmented and overlapping standards lead to inefficiencies and high compliance costs, which reduce adoption and effectiveness. A unified or harmonized standards framework, possibly tiered by application criticality, could enable scalable assurance while adopting commercial best practices.

Will emphasized that market dynamics are a major obstacle. Without clear demand signals, the industry lacks motivation to invest in reliable domestic capabilities. Therefore, assurance must be made economically feasible through standards, shared costs, and aligned incentives. Ultimately, a coordinated, industry-led, and government-supported approach is necessary to build resilient, trusted, and scalable electronics ecosystems.

2.1.6.1. Call to Action/Next Steps

- Develop a unified or tiered standards framework

  Consolidate fragmented standards into a coherent, interoperable framework aligned with allies. Reduce duplication and compliance burdens while enabling scalability across commercial and defense applications.

- Create strong demand signals for trusted electronics

  Government and industry must incentivize the adoption of assured components through procurement requirements, funding, and requirements to justify private-sector investment in domestic and allied supply chains.

- Prioritize high-risk components for assurance

  Focus efforts on critical, high-leverage components (e.g., processors, power management, memory, transceivers) for which supply chain risks and security impacts are greatest.

- Leverage commercial best practices and dual-use models

  Integrate proven industry practices into assurance frameworks and promote dual-use solutions to share costs, improve scalability, and accelerate adoption across sectors.

2.1.7. AI Provenance and Traceability

Jeremy Muldavin discussed how AI-driven design and semiconductor innovation are rapidly transforming the ability to build secure, traceable, and trustworthy systems. He highlighted the

convergence of “AI for design” and “design for AI,” where advanced chips enable AI systems while AI simultaneously accelerates chip development, verification, and system modeling. This evolution is critical as semiconductors underpin massive global markets, especially AI infrastructure, communications, automotive systems, and increasingly complex “physical AI” applications (e.g., robotics, autonomous vehicles systems).

A key challenge is building trust within the semiconductor and AI supply chains. Muldavin stressed that trust must start at the hardware level by ensuring IP integrity, validated firmware/software, and the secure management of AI models and data. He advocated for end-to-end traceability and provenance that track how chips are designed, manufactured, and deployed with verifiable inputs/outputs and non-repudiation. Without this, the large-scale adoption of AI, especially in critical infrastructure, will continue to be hindered by security concerns.

Muldavin suggested using AI-powered EDA tools and multi-agent systems to automate design, verification, and compliance processes, greatly increasing speed and coverage while integrating security requirements into workflows. These tools could facilitate IP validation, life cycle traceability, and automated security assurance at scale. Complementary innovations (e.g., chiplet architectures, digital twins, root-of-trust technologies) further support secure, modular, and scalable systems. However, challenges persist, including rising design costs, dominance by a few major players, fragmented supply chains, and a lack of economic incentives for secure, traceable systems. Muldavin emphasized the need for a global traceability framework that combines provenance, certification, analytics, and governance to enable both the physical and digital validation of chips throughout their life cycle. Ultimately, he argued that security and traceability must become measurable, valued, and integrated into the ecosystem to ensure trustworthy AI and resilient infrastructure.

2.1.7.1. Call to Action/Next Steps

- Establish global traceability and provenance standards

  Develop interoperable frameworks for tracking semiconductor IP, design processes, manufacturing, and deployment. Include physical and digital identifiers, certification mechanisms, and non-repudiation to enable verifiable trust across the life cycle.

- Embed security and compliance into AI-driven design flows

  Integrate security assertions, IP validation, and traceability directly into AI-enabled EDA tools and automated workflows to ensure “secure-by-design” outcomes at scale.

- Create economic incentives for secure supply chains

  Build market mechanisms and procurement preferences that reward traceability, provenance, and security, shifting from cost-driven decisions to value-driven trust in critical infrastructure.

- Drive public-private collaboration and pilot programs

Align industry, government, and academia to fund pilots, develop dual-use (e.g., commercial and defense) solutions, and accelerate the adoption of traceability frameworks, especially for critical infrastructure and national security systems.

2.1.8. Fundamentals of Semiconductor and Hardware Security for an AI World

Reed Hinkel emphasized that semiconductor and hardware security are vital for building trustworthy AI systems. He explained Synopsys' expanded capabilities after acquiring Ansys, highlighting a move toward integrated silicon-to-system design, simulation, and life cycle management. This integration enables more advanced AI-driven products, accelerates development timelines, and ensures that security is embedded throughout the design process.

A key theme was the need for widespread security across chips, systems, and data throughout their life cycles. As AI adoption increases, so do attack surfaces, including supply chain vulnerabilities and insider threats. Hinkel emphasized that security must be integrated from design to deployment, using technologies like cryptography, side-channel protections, and physically unclonable functions to ensure trust, provenance, and traceability. He highlighted the urgency of adopting post-quantum cryptography (PQC) and warned that current encryption methods may soon become obsolete. With regulatory deadlines approaching (e.g., 2030–2035 deprecation windows), organizations must act now, especially considering the long hardware life cycles in sectors like automotive and industrial systems. He also pointed out the risk of "harvest now, decrypt later" attacks, making immediate transition planning essential.

Hinkel also explored evolving standards and industry collaboration, particularly in cloud AI and chiplet-based architectures. Initiatives like the Open Compute Project and emerging interconnect standards increasingly incorporate security and attestation mechanisms. Finally, Hinkel discussed physical AI systems (e.g., robotics, autonomous systems), noting their complexity and the necessity of "secure-by-design" and high-assurance, fail-safe architectures to ensure safety and reliability in real-world environments.

2.1.8.1. Call to Action/Next Steps

- Accelerate post-quantum readiness

  Organizations must begin the immediate planning and deployment of PQC across systems and supply chains. Long life cycle devices should be prioritized to avoid future incompatibility, especially with impending regulatory deadlines and early industry shifts toward 2027 readiness.

- Embed security across the life cycle

  Adopt a "secure-by-design" approach that integrates security from initial chip design through deployment and operation. This includes implementing strong identity, provenance tracking, and continuous monitoring to mitigate risks across the supply chain.

- Strengthen supply chain trust and traceability

Develop mechanisms to verify component authenticity and track life cycle interactions. Techniques like hardware-based identities and the cryptographic validation of IP can help ensure integrity and reduce risks from tampering or counterfeit components.

- Engage in standards and industry collaboration

  Actively participate in emerging standards bodies and initiatives (e.g., chiplet security, Open Compute Project) to align with evolving best practices, influence secure architectures, and ensure interoperability in next-generation AI systems.

2.1.9. Siemens Tessent

Lee Harrison discussed how Siemens Tessent is evolving design-for-test (DFT) into a security-focused, life cycle-wide capability for modern semiconductors. He traced the evolution of DFT from early automatic test pattern generation and compression in the 2000s to more complex hierarchical designs and, more recently, Siemens' Streaming Scan Network (SSN), which enhances scalability and efficiency in testing large, modular chips.

As semiconductor designs become more complex, especially with AI-driven workloads, 3D ICs, and chiplet architectures, DFT infrastructure now offers deep internal access to nearly all nodes in a chip. While this allows for high-quality testing, it also introduces significant security risks. Harrison highlighted that traditional DFT, once limited to manufacturing, is now expanding into Silicon Lifecycle Management (SLM), where chips are continuously monitored in the field for performance, aging, failures, and potential cyber threats. SLM introduces new capabilities (e.g., in-field testing, functional monitoring, parametric telemetry) that generate vast amounts of sensitive data. This data enables predictive maintenance, failure diagnosis, and operational optimization, resulting in significant cost savings demonstrated through fewer automotive recalls and longer data center hardware lifespans. However, it also necessitates robust security for data access, transmission, and storage.

To address these risks, Harrison emphasized the importance of integrating root-of-trust architectures into DFT systems to secure interfaces, encrypt test data, and enable controlled, multi-level access for users and throughout the life cycle. The challenge becomes even more complex in multi-die and chiplet systems, where inconsistent security standards and diverse supply chains require interoperable and area-efficient solutions. Ultimately, security must be built into DFT and SLM as a fundamental element to ensure trust, scalability, and value.

2.1.9.1. Call to Action/Next Steps

- Embed security into DFT and SLM architectures by design

  Integrate root-of-trust, encryption, and authentication mechanisms directly into the test infrastructure to secure access to internal chip nodes and sensitive life cycle data from manufacturing through in-field operations.

- Develop standards for chiplet and multi-die security interoperability

Address cross-vendor security gaps by establishing common frameworks and standards to ensure consistent protection across heterogeneous supply chains and modular semiconductor designs.

- Leverage SLM data for predictive and economic value

  Invest in analytics pipelines that transform telemetry into actionable insights and enable predictive maintenance, targeted recalls, and extended hardware life cycles to reduce operational costs.

- Adopt multi-level access control and life cycle-aware security guidance

  Implement granular authentication schemes that adjust access privileges by user role and life cycle stage to ensure that only necessary features and data are exposed at each phase.

2.2. Government Panel

Christine Rink moderated a panel from a government/defense perspective that featured introductions from experts, such as Qiaoyan Yu (hardware security researcher focused on channels, embedded systems, and semiconductor life cycles), Jim Will (Partnership for Secure Electronics, with industry experience at SkyWater and DoD/NSA), Rob Aitken (Chips R&D, interested in chiplet security and infrastructure), Eric Eilertson (Microsoft Azure security architect, creator of OCP SAFE third-party audit framework, focused on hardware/firmware privacy against insiders and nation-states), and Robert

Watson (University of Cambridge professor leading CHERI project for hardware-software co-design and memory safety, plus CHERI Alliance director).

The panel centered on **defining microelectronics assurance** as evidence-based confidence in function and provenance across the semiconductor life cycle while acknowledging evolving threats in a “cat-and-mouse” game. Panelists emphasized that hardware security problems often intersect with software, requiring bidirectional information flow (e.g., secure-by-design principles, requirements from applications back to design). Key themes included realistic threat models (e.g., physical attacks like “break once, run anywhere” on encrypted memory, despite being declared out-of-scope by vendors), the composition problem in chiplets/SoCs, raising the security floor for ubiquitous low-end devices (e.g., IoT, Raspberry Pi-like platforms that end up in critical infrastructure), memory safety via technologies, the incentives and economics of adoption, and leveraging existing tools (e.g., static analysis, linting, audits) without boiling the ocean.

Panelists also stressed practical gaps, such as deploying known capabilities (e.g., linting for legacy code, supply chain verification), addressing low-cost physical attacks that democratize threats, standardizing definitions and practices (e.g., for memory safety to enable procurement and regulation), creating checklists and tiers of trust, and building consensus for incentives like third-party audits (e.g., OCP SAFE) that allow credit-claiming and reduce first-mover risk. They highlighted the need for living standards, industry-led best practices, shared reference designs, and metrics that demonstrate value to overcome commercial resistance to costs or regulation.

2.2.1. Call to Action/Next Steps

- Refine and standardize definitions of assurance, threat models, and memory safety practices

  Collaborate via organizations to create clear, evidence-based terms as well as checklists or tiers of trust to enable consistent frameworks, procurement signals, and adoption without ambiguity while accounting for application-specific risks.

- Prioritize raising the security floor for low-end and ubiquitous devices

  Focus research and standards on embedded, IoT, and microcontroller security (e.g., Raspberry Pi-like platforms); physical attack resistance; and debug infrastructure repurposing. Influence high-impact designs early, as they permeate critical infrastructure and data centers beyond high-end targets.

- Develop practical incentives and shared tools for broader adoption

  Promote simple and rewarding frameworks with third-party audits (e.g., expand OCP SAFE), industry-shared datasets for ML/research, reference designs, and metrics that demonstrate value. Combine these with demand signals from government and large buyers as well as living best practices to overcome economic barriers and first-mover risks.

- **Bridge hardware-software gaps through co-design and deployment interventions**

  Advance memory-safe interactions, static analysis, and linting for new and legacy code and secure-by-design principles with bidirectional requirements. Target vertical stacks (e.g., roots of trust, enclaves) for quicker wins while addressing composition in chiplets and SoCs.

3. Summary and Road Ahead

3.1. Establish Unified Standards and Governance for Hardware Security and Assurance

There is a strong consensus on the need to reduce fragmentation by establishing shared, evidence-based frameworks across assurance, threat modeling, memory safety, provenance, and life cycle security. This includes developing tiered trust models, standardized terminology, and harmonized guidance aligned across international organizations and allied ecosystems. A coordinated governance structure with clear stakeholder roles (e.g., industry, government, academia) should outline a multi-year plan with measurable goals. Standards should cover the entire stack from silicon to systems and include emerging areas, such as chiplets, AI hardware, and post-quantum readiness. Consolidation will enhance interoperability, reduce compliance burdens, and accelerate adoption.

3.2. Embed Security, Provenance, and Traceability Across the Full Semiconductor Life Cycle

A recurring theme is the need to make security intrinsic rather than additive. This includes integrating secure-by-design principles into EDA/AI-driven workflows, embedding hardware roots of trust, and incorporating security checks into design-for-test (DFT), supply chain processes, and in-field operations. Provenance and traceability mechanisms that cover IP, components, manufacturing steps, and system deployment should leverage cryptographic identities, SBOMs, attestation, and verification pipelines. Interoperable global traceability standards and life cycle-aware access controls will enable trusted provenance, reduce counterfeit risks, and support compliance, auditability, and incident response across complex, multi-vendor ecosystems.

3.3. Strengthen Supply Chain Security and Economic Incentives

Adoption barriers are mainly economic and structural. Stakeholders stressed the importance of aligning incentives through procurement requirements, government demand signals, and industry consortia programs to reward secure components and verified supply chains. Requirements (e.g., SBOMs, security clauses in contracts, third-party audits) help establish consistent expectations, and focusing on high-risk, high-impact components (e.g., processors, memory, power systems, transceivers) maximizes effectiveness. Public-private partnerships, pilot programs, and dual-use (i.e., commercial and defense) approaches can reduce investment risks while demonstrating tangible returns for secure-by-design and provenance-enabled ecosystems.

3.4. Advance Scalable Tools, Automation, and Validation Methods

To make assurance practical at scale, investments are necessary in automation and tooling, including AI-assisted analysis, static verification, fuzzing, and formal methods designed for hardware and firmware. Standardized testing and validation methods — especially for AI chips, side-channel resistance, and fault injection resilience — should be developed and aligned with

evolving threat models, including both deterministic and adaptive attack techniques. Shared datasets, reference designs, and benchmarking frameworks can help accelerate research and evaluation. Incorporating security into automated workflows and using telemetry for predictive insights will enable continuous assurance, monitoring, and life cycle management.

3.5. Build Workforce Capacity and Collaborative Ecosystems

A vital cross-cutting requirement is to build a skilled workforce and promote ongoing collaboration across sectors to sustain long-term progress. This involves connecting academia and industry through hands-on training, turning institutional knowledge into accessible resources, and increasing participation in NSF-backed, CHIPS R&D, and interagency programs. Participating in workshops, consortia (e.g., OCP, ASCET-like efforts), and standards organizations will help align research with practical needs. Fostering interdisciplinary teams and incentivizing knowledge sharing, open datasets, and community-driven frameworks will help build expertise, break down silos, and ensure that advances in hardware security are widely adopted and continuously refined.

Appendix A. Workshop Agenda

Start	Topic	Who
8:30	NIST Safety Briefing	Jay Rekhi, NIST
8:35	Welcome and Introduction to NIST	Jim St. Pierre, NIST
8:50	Opening Remarks by Organizing Committee	Jeyavijayan Rajendran (Texas A&M), Gang Qu (UMD), **Ahmad-Reza Sadeghi (TU Darmstadt)**
Session 1: Chair Ahmad-Reza Sadeghi		
9:00	A Roadmap Towards Scalable Security Assurance	Jason Fung, Intel
9:30	Security in Memory and Storage	Tamara Schmitz, Micron
10:00	Hardware Attacks on AI Inference Chips	Marc Witteman, Keysight
10:30	**Break**	
Session 2: Chair Jeyavijayan Rajendran, Texas A&M		
10:45	NIST Semiconductor Standards Activities	Mary Bedner, NIST
11:15	NSF/CISE Security, Privacy, and Trust in Cyberspace (SaTC 2.0)	Selcuk Uluagac, NSF
11:45	Standards-Based Assurance for Trusted Electronics	Jim Will, USPAE
12:15	**Lunch**	
Session 3: Jay Rekhi, NIST		
13:00	AI Provenance and Traceability	Jeremy Muldavin, Cadence
13:30	The Fundamentals of Semiconductor & Hardware Security for an AI World	Reed Hinkel, Synopsys
14:00	Siemens Tessent	Lee Harrison, Siemens
14:30	**Break**	
14:45	**Government Panel** Moderator: Christine Rink, MITRE Panelists: Qiaoyan Yu (NSF), Jim Will (USPAE), Robert Aitken (NIST CHIPS), Eric Eilertson (Microsoft), Robert Watson (University of Cambridge)	
15:45 16:45	**Breakout Sessions** Challenges in adopting hardware security in semiconductors, Moderators: Cvitan, Sandra (NIST CHIPS) Secure Lifecycle Management, Moderator: Matthew Areno (Rickert-Areno Engineering, LLC) Emerging Threats, Moderator: Jason Fung (Intel) Breakout Session Summary Report	
17:00	**Closing Remarks**	

Fig. 1. SUSHI@NIST workshop agenda