NIST Interagency Report
NIST IR 8532 ipd

Workshop on Enhancing Security of Devices and Components Across the Supply Chain

Initial Public Draft

Sanjay (Jay) Rekhi
D. Richard Kuhn
Kim Schaffer
Murugiah Souppaya
A.J. Stein
Noah Waller
Nelson Hastings
Michael Ogata
William C. Barker

This publication is available free of charge from:
https://doi.org/10.6028/NIST.IR.8532.ipd

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

NIST Interagency Report
NIST IR 8532 ipd

Workshop on Enhancing Security of Devices and Components Across the Supply Chain

Initial Public Draft

Sanjay (Jay) Rekhi
D. Richard Kuhn
Kim Schaffer
Murugiah Souppaya
A.J. Stein
Noah Waller
Computer Security Division
Information Technology Laboratory

Nelson Hastings
Michael Ogata
Applied Cybersecurity Division
Information Technology Laboratory

William C. Barker
Dakota Consulting

This publication is available free of charge from:
https://doi.org/10.6028/NIST.IR.8532.ipd

August 2024

U.S. Department of Commerce
Gina M. Raimondo, Secretary

National Institute of Standards and Technology
Laurie E. Locascio, 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

Publication History

Approved by the NIST Editorial Review Board on YYYY-MM-DD [Will be added to final publication.]

How to Cite this NIST Technical Series Publication

Rekhi S, Kuhn DR, Schaffer K, Souppaya M, Stein AJ, Waller N, Hastings N, Ogata M, Barker WC (2024) Workshop on Enhancing Security of Devices and Components Across the Supply Chain. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Interagency Report (IR) NIST IR 8532 ipd. https://doi.org/10.6028/NIST.IR.8532.ipd

Author ORCID iDs

William C. Barker: 0000-0002-4113-8861
Nelson Hastings: 0000-0003-2444-6413
D. Richard Kuhn: 0000-0003-0050-1596
Michael Ogata: 0000-0002-8457-2430
Sanjay (Jay) Rekhi: 0009-0008-8711-4030
Kim Schaffer: 0000-0003-3073-2395
Murugiah Souppaya: 0000-0002-8055-8527
A.J. Stein: 0000-0003-1092-2642
Noah Waller: 0000-0002-6979-9725

Public Comment Period

August 14, 2024 – September 16, 2024

Submit Comments

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/8532/ipd, including related content, potential updates, and document history.

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

Abstract

NIST hosted an in-person, all-day workshop on February 27, 2024, to discuss existing and emerging cybersecurity threats and mitigation techniques for semiconductors throughout their life cycle. The workshop obtained valuable feedback from industry, academia, and government to inform NIST's development of cybersecurity and supply chain standards, guidance, and recommended practices. The discussion focused on semiconductor development and highlighted cybersecurity measurements and metrics that utilize reference data sets to facilitate the testing, attestation, certification, verification, and validation of semiconductor components. It also emphasized the use of automated cybersecurity tools and techniques to secure manufacturing environments throughout the development life cycle. This report summarizes the content that was presented and discussed at the workshop.

Keywords

cybersecurity; hardware security; measurement; semiconductors; supply chain; vulnerabilities.

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

List of Figures

1. Introduction

Semiconductor-based hardware is the foundation of modern-day electronics — from smartphones, computers, and telecommunications to transportation and critical infrastructure. The semiconductor hardware supply chain is a complex network of companies that collectively provide intellectual property, designs, and raw materials and manufacture, test, package, and distribute products. Coordination among components of a supply chain is required at different stages, including inception, deployment to end users, maintenance during use, and disposal or end-of-life. Securing semiconductor-based hardware and their supply chains helps protect sensitive information, maintain the integrity of systems, and ensure overall stability across the infrastructure and connected world.

Securing semiconductors involves the security of the component being built as well as the design, development, manufacturing, and distribution environments. Figure 1 illustrates the components of achieving robust semiconductor security.

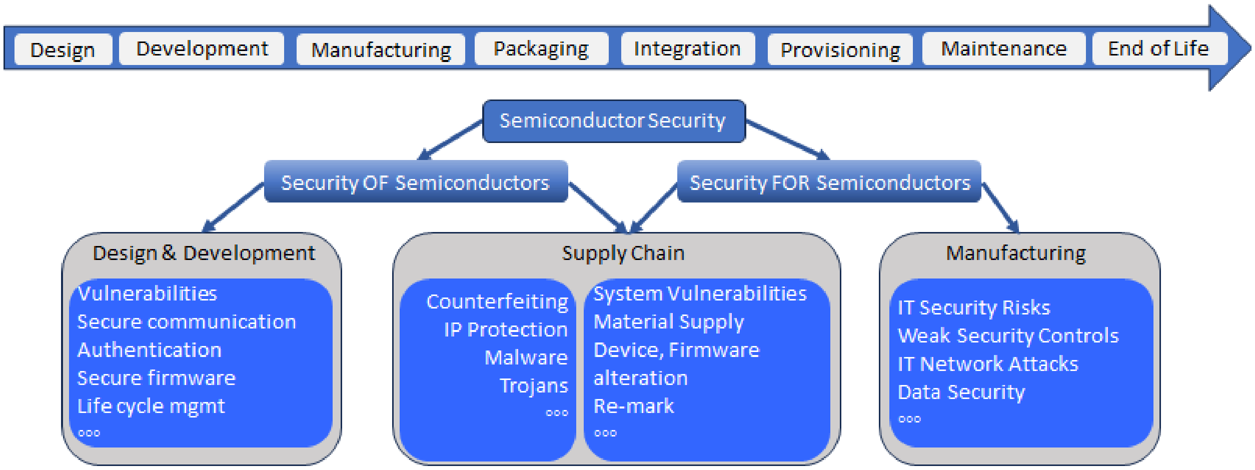

Fig. 1. Components for securing microelectronics

These activities range from physical protection mechanisms (e.g., tamper-resistant packaging) to strong encryption protocols to safeguard data. Secure boot processes ensure that only verified firmware runs, and maintaining supply chain integrity through verification and audits prevents counterfeit components. Life cycle management includes secure provisioning, updates, and end-of-life processes that are complemented by rigorous security testing and compliance with regulatory standards. Educating users on secure practices and continuously improving security measures further fortifies defenses against evolving threats and ensures that semiconductor devices operate securely throughout their life cycle.

The National Strategy on Microelectronics Research has emphasized the prioritization of hardware integrity and security. In response, NIST convened its inaugural workshop, “Enhancing Security of Devices and Components Across the Supply Chain”[1] on February 27, 2024, at the National Cybersecurity Center of Excellence (NCCoE) facility. At this workshop,

[1] See https://csrc.nist.gov/Events/2024/enhancing-security-of-devices-and-components.

government, academia, and industry experts gathered to collaborate on research efforts, drive innovation, and establish standards, guidance, and practical implementations in a rapidly evolving landscape. The workshop was primarily an in-person event with a few remote speakers and participants. In total, there were 98 participants, with 79 participants almost evenly distributed from government and industry. The remaining 19 attendees were from academia and standards development organizations (SDOs). Figure 2 shows the distribution of participants.

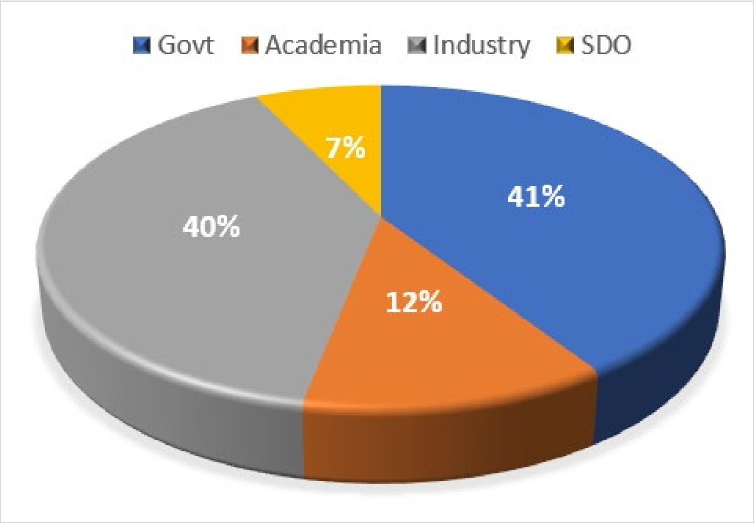

Fig. 2. Distribution of workshop participants

2. Workshop Sessions

2.1. Hardware Development Life Cycle

Semiconductor and integrated circuit (IC) development is complex, non-linear, and varies from one manufacturer to another. The specialization of technologies required to design modern-day hardware has democratized the life cycle across national boundaries. This panel discussion focused on open security concerns that face the hardware development life cycle.

2.1.1. Speaker Viewpoints

The panel opened with Jonathan Ring, Deputy Assistant National Cyber Director for Technology Security from the Office of the National Cyber Director (ONCD). Jonathan reflected that ONCD's National Cybersecurity Strategy recognizes the importance of improving the cybersecurity of the Nation's critical infrastructure, which includes restoring the production of critical goods to the United States (US) as well as improving cybersecurity in the semiconductor/IC supply chain.

Jonathan highlighted past achievements of the Biden-Harris Administration that overlap with cybersecurity in the hardware development life cycle, including:

- The notion of shifting the balance of security to those best suited to bear it, as represented in the Cybersecurity & Infrastructure Security Agency's (CISA) Secure by Design initiative
- The NIST CHIPS R&D Metrology Program, which outlines gaps in the semiconductor ecosystem
- Ongoing work through the Subcommittee for Microelectronics Leadership
- The focus on eliminating entire classes of vulnerabilities in software in their report, Back to the Building Blocks: A Path Toward Secure and Measurable Software

Furthermore, Jonathan stressed the importance and continued need for public-private partnerships, like the Semiconductor Research Corporation and the Semiconductor Industry Association. He also stressed the need for continued conversations concerning advanced metrology for supply chain trust and assurance, guidelines for security analytics and automation, vulnerability management across all product life cycles, and the use of machine learning in chip design and manufacturing.

The second speaker for the panel was Adam Golodner, advisor to the Semiconductor Industry Association. Adam spoke about the importance of leveraging and adapting NIST resources, like the Cybersecurity Framework (CSF). Adam highlighted four key reasons why a framework for hardware security and supply chain that is similar to the CSF would be useful to the semiconductor industry:

1. Frameworks like the CSF enable thoughtful, flexible, and configurable approaches to security and allow enterprises to adopt the processes and maturity level that best move security forward for them.

2. A hardware security and supply chain framework can benefit from the recognition and adoption that the CSF has already earned.
3. The CSF and NIST have global reach and acceptance. Similarly, establishing an internationally recognized set of best practices will enhance security and innovation in a global hardware and supply chain environment.
4. NIST has the pedigree for getting security "right."

Furthermore, Adam reflected that while security is core to many company brands, quantifying its return on investment is a C-suite and Board of Directors issue.

The last panel member was Matthew Areno, Senior Principal Engineer for Intel Corporation. Matthew admitted that there is no industry standard to describe the phases of the supply chain for ICs. However, his expertise and history with the industry have enabled him to develop his own conceptual model that is divided into seven stages: concept, development, integration, manufacturing, test, provisioning, and deployment. During his time at Intel, Matthew has worked to develop a threat model for Intel's supply chain. Intel discovered that each of their teams approached the exercise of threat modeling in different ways, which made it difficult and time-consuming for them to exchange information. As a result, Intel has been developing and deploying a unified Threat Modeling Tool that utilizes vulnerability resources like MITRE Common Weakness Enumeration (CWE) to automatically make design-based suggestions on potential vulnerabilities and mitigations. The tool also forms a closed loop that incorporates knowledge learned from Intel's own design processes.

Matthew also described Intel's Transparent Supply Chain Initiative, whose goal is to provide customers with provenance and integrity data for various components during the supply chain process. He highlighted Project Amber — Intel's Trust Authority for validating execution environments.

2.1.2. Key Highlights

The following key points were identified during the speakers' presentations and the Q&A portion:

- There is a definite need for continued public/private partnerships.
- There is a need for more standards around how the supply chain process is structured.
- There is a need for more standardized methods for threat modeling.
- Quantifying the return on investment for security in the supply chain is still difficult, making it challenging for those who observe deficiencies to justify addressing them to higher-level management. Possible avenues of advancement include:
    - Developing more standardized security metrics
    - Continuing to leverage and develop CSF-style resources to better communicate risk
    - Framing security as a service that can be used as a revenue stream

- There is a tension between hardware security and software security due to the vastly different cost in remediating vulnerabilities at the hardware level.

2.2. Metrology

This panel focused on open issues and concerns related to security metrology and metrics (e.g., metrics for design-for-trust techniques, metrology for the holistic assessment of power side-channel leakage across the development life cycle, metrology for analog signal security) for semiconductors and ICs throughout the hardware development life cycle.

2.2.1. Speaker Viewpoints

This session began with Lok Yan from DARPA's Microsystems Technology Office (MTO), whose portfolio includes the Automatic Implementation of Secure Silicon (AISS) program. Lok provided a high-level overview of why the semiconductor community needs to recognize security metrics. Meaningful metrics require assets, use cases, and threats to be clearly identified within the context of a threat model. The value of assets, associated potential threats, and the consequence of an asset's compromise must be quantifiable to support the design decision-making process.

Understanding and identifying a minimum acceptable security baseline within the context of a particular use case would help bring various security metrics together. Due to the continuously evolving threat space, Lok emphasized that establishing an initial set of metrics and an associated minimum security baseline is not a one-time process. Rather, they must be continuously measured, monitored, and updated as assets, threats, vulnerabilities, and use cases evolve over time. Finally, security metrics could support decisions related to the threat and vulnerability space versus the time to test or implement countermeasures. More time must be allocated to implement and test countermeasures that address the most important threats and vulnerabilities for a particular use case.

Jason Oberg, Cycuity CTO and co-founder, described how metrics can be used in the hardware security space based on his organization's experience. Since different metrics are relevant to different people within an organization (e.g., metrics that work for a designer or tester may not translate to executives), there need to be different tiers of metrics based on the audience or user. For example, the CWE structure could be used to define security requirements that lead to associated security metrics. Jason noted that CWEs for hardware are relatively new, starting around 2020, compared to CWEs' long track record for software, which started around 2006. Since CWEs point out common root cause weaknesses that lead to vulnerabilities, they foster a more proactive approach by allowing security issues to be discovered and mitigated earlier in the hardware development life cycle, thus lowering the financial impact. Having good metrics that can be used to inform design and business decisions would be helpful to the semiconductor community.

Finally, Mark Tehranipoor from the University of Florida emphasized the importance of thinking about security and associated metrics in the early stages of the hardware development life cycle. Specifically, this involves conducting a security and risk assessment and developing a security architecture during the specification and planning phases. The dynamic nature of the IC development process may impact the metrics that need to be developed, and a security metric that is suitable for the register-transfer level (RTL) may not be appropriate or work at the gate or physical level. He noted that time-to-market constraints still run on an approximately six- to nine-month cycle, but complexity has increased, leading to the need for automation in security, reliability, and testing. This increased complexity also increases the number of assets to protect and the number of potential vulnerabilities to mitigate. Regardless of what security techniques and associated security metrics are used early in the development life cycle (i.e., pre-silicon), verification is still needed during the physical layout and post-silicon stages. Finally, he identified the need for the continued development of security solutions at the material, physical, and device levels of semiconductors.

2.2.2. Key Highlights

The following key points were identified during the speakers' presentations and the Q&A portion:

- For security metrics to be meaningful, they need to be provided in the context of threat models, use cases, and vulnerable assets.
- Security metrics will need to be tailored based on where they are in the development life cycle (e.g., functional versus physical design stages) and to whom they are communicated (e.g., design engineer versus executive).
- Security must be on par with other design constraints (e.g., area, power, performance, and reliability), so having a community agreed-upon minimum security baseline might be a good first step.
- There is benefit in enhancing design tools to support security techniques and practices via automation.
- There is an opportunity to investigate the potential application of software-based security techniques and practices within the hardware domain.

2.3. Hardware/Silicon Testing

Speakers from Synopsys, PQShield, and the University of Maryland shared their expertise and vision of where the industry and technology are headed.

2.3.1. Speaker Viewpoints

Mike Borza from Synopsys presented "Security Verification of SoC Hardware," an overview of the current status, new developments, and likely future progress in ensuring security for system-on-chip (SoC) designs. Security has begun to drive the design requirements of SoCs, which has resulted in tool vendors adding features for strong verification. Interoperability needs are also driving work on standards, such as IEEE P3164, Security Annotation for Electronic Design Interchange. Additionally, security requirements now feed into every aspect of the architectural specification of a new chip and subsequent RTL design and analysis. Designs must include features such as secure boot, secure memory and interfaces, and hardware countermeasures to meet the goals of avoiding vulnerabilities. Static and dynamic verification approaches involve concentrating on formal methods early in the design process, which is then followed by simulation and testing that meet stringent coverage criteria. Mike also described tool support for all aspects of a security-focused verification platform, including the regular functional verification of security functions, checking on-chip data propagation to secure data at rest or in motion, and investigating possible tampering or intrusion.

It is anticipated that tool support will be improved to guarantee high levels of coverage with an improved ability to reason about the physical realizations of designs rather than only abstract descriptions. These improvements will be enhanced by developing standards to describe and communicate security information on designs and may eventually benefit from artificial intelligence/machine learning systems that incorporate knowledge of threats and potential weaknesses.

Niels Samwel of PQShield described his company's work on "Automation for Side-Channel and Security Testing of Hardware IP." Cybersecurity testing services provided by PQShield include side-channel testing and quality assurance of hardware designs. Side-channel tests can be conducted according to Federal Information Processing Standard (FIPS) 140-3 guidelines for cryptographic hardware and using Common Criteria recommendations. Using FIPS 140-3, it is possible to provide remote operation with a sufficient quantity of tests to allow for statistically valid pass-fail results. Common Criteria side-channel testing capabilities are included to estimate the number of traces required for key recovery. These test methods also make it possible to target specific vulnerabilities and attack types, including template attacks, key recovery attacks, and both correlation power analysis and differential power analysis.

Product quality testing services are offered for multiple phases of hardware and software product development. For digital circuit design, linting and automated field-programmable gate array (FPGA) functional testing and design implementation evaluations are provided with verification phase capabilities that include constrained random verification, coverage measures, and bounded model checking. Software assurance capabilities include static analysis and the automated testing of implementations. Verification phase processes include unit, integration, and system-level tests that also measure test coverage.

PQShield has integrated these testing services to develop a three-level scale for security that is tied to the levels defined in FIPS 140-3 and the Common Criteria:

- The Cloud Level of the PQShield scale targets safety against fuzzing and remote attacks and corresponds to FIPS 140-3 Level 1 or CC EAL1 and AVA_VAN.1.
- The Edge Level of PQShield evaluates safety against "push button" physical attacks and corresponds to FIPS 140-3 software Level 2 and hardware Level 3 or CC EAL2 to 3 and AVA_VAN.2.
- The Government Level is the highest level of the PQShield scale for safety against expert labs and corresponds to FIPS 140-3 software Level 2 and hardware Level 4 or CC EAL4+ to 7 and AVA_VAN.5.

The security-level scale and associated tests are intended to allow organizations to select cybersecurity evaluations according to their risk management needs.

Ankur Srivastava of the University of Maryland (UMD) College Park presented research on "Verification and Validation for Hardware Security Constructs," which focuses on design obfuscation and trojan detection and mitigation measures. The need for design obfuscation arises from current practices in which a fabless IC designer outsources production of the chip to an offshore foundry. The potential risk of outsourced intellectual property piracy or counterfeiting affects both defense and industry customers, as well as the company that created the design. Logic locking obfuscation techniques have been developed to mitigate these risks, but sound measures of resistance to attack are also needed. The most researched scenario for evaluating obfuscation resistance is the case in which an attacker has a working chip that enables them to infer a design from input-output pairs or sophisticated imaging. There is less research for cases in which attackers do not have information on the design or have only a library of similar designs. Researchers have developed an extensive set of techniques to identify potential weaknesses to zero-knowledge or partial prior knowledge of designs.

Hardware trojans are another source of concern in hardware security. The potential exists for a malicious function to be included in a chip and triggered later by an attacker who knows the key that can be included in inputs. UMD researchers are investigating vulnerability and detectability analysis for trojan mitigation schemes. This work includes statistical analysis to determine trojan triggers and a large study evaluating the trade-offs between the likelihood of detection and the rarity/complexity of the trojan trigger using measures such as trigger length or the size of a finite state machine space that must be traversed to initiate the malicious function. This is accomplished by stress testing a spectrum of trojan types that are implemented for evaluation purposes. The area, power, and performance overheads of trojans must be evaluated because of the limits on detecting trojans by testing. Ankur emphasized the value of a strategic, layered approach to vulnerability analysis and the need for sound mathematical models (e.g., those separating trojans from bugs) to consider an attacker's different levels of access, knowledge, and control. UMD is also developing sound metrics for hardware security constructs and security strategies for heterogeneous integration.

2.3.2. Key Highlights

The session speakers identified several needs and near-term expectations. A common theme was the need for an integrated approach to hardware security that includes advanced capabilities for all aspects of the problem. In particular, the industry should be focusing on:

- Better tool support to ensure more complete design coverage. Tool advances should also include formal approaches to reason about the physical realizations of designs beyond the current formal methods that focus on abstract representations and hardware description languages (HDLs).

- More standardized interoperability of tools and input/output. Currently, semiconductor companies tend to have their own collections of specialized tools, which makes it difficult to share information with others in the industry.
- Improved data collection and understanding of vulnerabilities. This will allow for better risk management that aligns organizational risk tolerance with appropriate levels of analysis and testing. FIPS 140-3 and the Common Criteria are useful for analyzing test and assurance approaches for deterring particular attack classes and vulnerabilities.
- Better design obfuscation techniques, as well as vulnerability and detectability analysis of malicious insertions in fabricated designs. Among the most significant risks in today's offshoring environment are the loss of intellectual property and the potential for adversaries to compromise chips with hardware trojans. Given the limitations of detecting such vulnerabilities through testing alone, hardware analyses that include power and performance overheads are essential for identifying trojans and other chip malware.

2.4. Vulnerability Management

Hardware vulnerability management shares similar challenges with well-established software vulnerability management practices and also faces its own unique challenges. The panel discussion presented three perspectives around this theme: Qualcomm's present-day experience performing vulnerability management at scale, a Battelle researcher's futurist view of defending and attacking hardware with generative AI techniques, and NIST's view on the past and present of bug classification as it applies to hardware.

2.4.1. Speaker Viewpoints

Dan O'Loughlin described how the security work of the architecture, engineering, and evaluation teams for Qualcomm's SoC portfolio drives their vulnerability management program. As he noted, Qualcomm suggested doing this at scale, thinking holistically about security and vulnerability management. Vulnerability management is an integral part of Qualcomm's overall security assurance process, for which they suggested maximizing the best outcome for planned investment. Dan recommended the categorization of the root causes for vulnerabilities throughout the life cycle (e.g., pre- and post-silicon) and how to feed back into ongoing investments and operations. The most common cause is process compliance failures, while the second most common is specification traceability gaps. Therefore, their automation has focused on addressing these causes.

Qualcomm suggested focusing on countermeasures for missing threat assessments, which is an important root cause. Dan's team has made additional efforts in generating and maintaining automated threat models with the help of machine-readable data formats, such as SysML. This focus on the threat model, test plan, and supporting automation allows them to scale important security checks throughout the life cycle with available staff. It is important to match threat assessment and automated testing with vulnerability detection and analysis early in the life cycle. Qualcomm recommended investing heavily in this detection with fault injection, side-channel, and other techniques for pre-silicon testing. As Dan explained, this shift-left strategy is especially important for their products to detect and prevent vulnerabilities as early as possible before final certification and release to market. With all of this internal security evaluation and validation, Dan and his team have measured vulnerabilities and countermeasures over multiple generations of SoCs to confirm that the severity of findings is trending down over time.

However, publicly disclosing more detailed data is a different matter. Dan has been closely following regulatory changes for responsible disclosure in the software industry, but hardware vulnerability disclosure is fundamentally different. In his view, regulators and manufacturers have very different incentives, so partnership and further discussion will be required.

Next, Jeremy Bellay talked about Battelle's research regarding the impact of context on proper vulnerability management. Vulnerability categorization, like CWE, provides valuable foundational context. However, higher-level, human-friendly context is still resource-intensive to produce and error-prone. One such example is the reachability analysis — how accessible a target system with a given vulnerability is to an attacker. Another example is attack chain design, where each vulnerability disclosure provides attackers more opportunities to combine multiple vulnerabilities to fully exploit a system.

In the past, it has been difficult to organize data for higher-level contextual information using datasets and standards from NIST, MITRE, and others. However, Jeremy's team has recently utilized generative AI tooling to obtain this higher-level context without the additional resources needed for conventional methods. For him, the emergence of AI tools has moved the industry from the "age of context" to, as he terms it, the "age of interface." With this perspective and tools at hand, Bellay presented his success with advanced generative AI to augment the development of attack chains with vulnerability information. This approach shows promise, yet it is not devoid of risks. Jeremy presented examples of using generative AI systems with prompts defining strict policies that tools violated, despite being given the needed context. Nonetheless, he is confident that they will improve in this age of interface and enable new capabilities for attacker and defender alike.

Finally, Peter Mell presented his research on the software, hardware, and trends for vulnerability management in the past, present, and future. Historically, claims of unbreakable secure software were met with skepticism, while hardware was perceived to be the immutable root of trust. This perception persisted, even though hardware is designed and programmable with software. As Peter put it, in some sense, "hardware is software."

To effectively compare and contrast hardware and software, more data is needed for hardware vulnerability research. Peter compared the public infrastructure for software (e.g., CVE, NVD, CVSS, EPSS, KEV, CWE) to the current hardware vulnerability landscape. For the categorization of vulnerability types (CWE), less than half have been observed with confirmed hardware bugs, with little overlap in categories between hardware and software bugs. Additionally, Peter pointed out that he found little public evidence of hardware bugs, as opposed to thousands a year for software. He concedes that there are still some differences between hardware, software, and their vulnerabilities. Nonetheless, the paucity of data demonstrates room for improvement and a challenge to the hardware industry as it matures vulnerability management practices. This research did not uncover any obstacle to utilizing public software vulnerability infrastructure for hardware vulnerability management, and he welcomes work in this area.

At the end of the panel, attendees asked questions about the tools and processes that aid in traceability and security. Dan described a variety of tools for security and traceability management in existing projects and greenfield projects. He repeated his praise for model-based systems engineering and tools. Jeremy agreed that those tools were helpful. There were also questions regarding the presenters' views and techniques for resourcing the experts and funding needed for outcome-focused vulnerability management. Dan explained that Qualcomm prefers actuarial methods. Peter and Jeremy emphasized the importance of methods for higher-level context and how to encourage more research in that area to support outcome-focused vulnerability management.

2.4.2. Key Highlights

The following are some key takeaways for future work based on the presentations, feedback, and questions from the audience:

- Hardware vulnerability management could leverage public software vulnerability infrastructure, but more dialogue is needed to understand how to use it to meet hardware vendors' needs.
- Automation-friendly traceability techniques are necessary to scale the prevention and detection of vulnerabilities.
- Model-based systems engineering tools, techniques, and standards are in use, but it still needs to be determined how to measurably expand their use and incentivize industry.
- More comprehensive vulnerability data is needed to improve the taxonomies of hardware bugs that are understood and to predict those that are not.

2.5. Standards

This session explored various aspects of hardware security standards related to semiconductor manufacturing.

2.5.1. Speaker Viewpoints

Jeremy Muldavin represented the SAE G-32 committee and was the first presenter. He began by pointing out that while there are incentives for CHIPS fabricators to build in the US, there is currently very little market preference for an assured supply. If not corrected, the market will revert to focusing on buying cheaply, which will lead to a loss of investment. By creating standards that integrate assurance through traceability and provenance into systems engineering, the market can understand, measure, and adapt to the demand for a long-term assured supply. Jeremy stated that when the US promoted "Buy US, Build US," European customers were interested in a US supply chain, but when there appeared to be no teeth behind it, interest was lost.

With the amount of R&D invested in the National Semiconductor Technology Center (NSTC) and similar efforts, value must be added through assurance, or else cheaper products will take over the market again. Without knowing how to build programs that are directly tied to semiconductor manufacturing and show measurable assurance, research investments are being wasted. Agreed-upon measurable assurance requires believable standards that illuminate supply chains, identify market risks, create a basis for monetizing supply and security, and measure the impacts of assured supply. They also need to identify methods for immutable physical traceability, validate roots of trust, and identify ways to develop consumer-level traceability tools. The assurance is needed early, when the chips are inexpensive. The payoff is at the product and services end, where the applications have a far greater revenue stream.

Semiconductor manufacturing harvests a significant amount of data. Jeremy stated that Global Foundries accumulated about 12 terabytes of data per day. There must be an analytic environment to take advantage of that data through an "observe, orient, decide, and act" (OODA) loop to develop assurance and awareness capabilities (e.g., supply chain and digital twins capabilities) and support the stress testing of manufacturing supply disruptions in a way similar to the 2008 financial crisis bank stress testing. This would enable the transition from a trust framework of people watching people make assessments to using digital twins to model sensors and data to monitor the supply chain in a manner that creates value by establishing provenance and traceability.

Andrew Seward introduced the Semiconductor Manufacturing Cybersecurity Consortium (SMCC) efforts that Semiconductor Equipment and Materials International (SEMI) is implementing. SEMI is an international organization that has focused on the semiconductor industry since 1970 and provides global advocacy and technical leadership. It currently meets the cybersecurity needs of the industry, from material and equipment suppliers to end users. In November 2023, the SMCC and NIST met to identify key areas and seek volunteers for both internal leadership and action. Those attending the meeting made a good representation of all semiconductor-related industries. Starting with about 40 in-person and almost 70 online attendees, the volunteers from that group have grown to about 50+ since January 2024 and have support from several CSOs from major organizations.

Jennifer Lynn continued the SEMI presentation and stated that the SMCC is gaining momentum. The two days of whiteboard sessions at the November 2023 meeting established seven working groups: 1) factory cybersecurity implementation, 2) compliance readiness, 3) supply chain cybersecurity, 4) regulation and other specs, 5) threat sharing, 6) cybersecurity pre-standards engineering, and 7) outreach. Jennifer is leading working group 4, which will co-author the industry profile to map SEMI requirements to CSF 2.0. Anyone who wishes to aid in these efforts is encouraged to email cybersecurity@semi.org to talk to the working group leads and discuss how you might help. This work will establish the requirements for moving forward as well as how the existing structure can be protected for the rest of its lifetime.

The SAE G-32 is working on integrating cybersecurity assurance into a CPS systems engineering and product-focused process. The SEMI SMCC will transition the design and manufacturing floor to one that incorporates auditable cybersecurity.

A comment from the audience suggested that the SEMI standards for traceability at the wafer level and IPC's standards on traceability for manufacturing are a good basis for forming a liaison activity between SEMI, IPC, and SAE. NIST and other SDOs (e.g., IEEE) and entities (e.g., IT-ISAC) could benefit from coordinating and participating in such efforts.

Another comment asked whether NIST would consider leading the effort to assess the approximately 300 related standards that already exist, potentially by employing graph analytics and other AI searching, learning, and parsing tools. It was agreed that a unified view of cybersecurity standards would help many entities better understand what is available for use and what is needed, especially with the participation of the semiconductor-related corporations that are working with the CHIPS program.

This opened related discussions on the use of cybersecurity standards, such as how one decides which standards to apply based on a product or organization. A comment was also made about approaching the C-suite for a semiconductor business, which appears to be more interested in a semiconductor-focused cybersecurity standard than a generic cybersecurity standard. Another comment noted that creating a new standard for cybersecurity should leverage work that has been done in other areas, such as automotive or health care, and confirm applicability rather than "reinventing the wheel." In addition, requirements in standards need to be measurable. This can be difficult, as requirements are often created separately from the compliance aspect, and finding the right balance can be challenging.

A related comment focused on harmonizing the measurements in standards into a common or related set of metrics. In the future, it will be desirable to tailor requirements by referencing applicable parts of standards. The customer will need to assess their demand and available supply in order to verify the level of assurance required for products and services to meet their needs.

2.5.2. Key Highlights

While IT standards have continued to mature, awareness of the need for hardware standards and the importance of supply chain assurance, manufacturing policies, and resilience has only begun to grow. Cost cannot be the only consideration for semiconductor manufacturing; security and assurance value propositions and end user demand must be considered as well.

In the wider scope of current international supply chains, the integration of security and assurance measures is primarily relegated to larger manufacturers. Additional measures need to be uniformly integrated for both semiconductor manufacturers and their suppliers, from sophisticated equipment to raw materials. SEMI International has initiated an effort to gather all manufacturers and security and assurance experts to develop a set of standards that can be integrated into all businesses across the supply chain, with verification being a major component. This effort will reference existing IT and other non-semiconductor industry standards and work with other SDOs when such standards do not exist.

2.6. Closing Remarks

Serge Leef, the Secure Microelectronics Design, Implementation, and Fabrication Enablement Lead at Microsoft Azure, provided closing remarks. In "Challenges and Opportunities in Commercializing Security Research," he addressed market barriers for hardware security products and provided an overview of the market segments:

1. Huge/large organizations for whom hardware security is a critical need and that have large teams of experts who develop appropriate solutions to address their needs across multiple high-value, large-scale products
2. Mid-size semiconductor and system companies that understand the need but lack the expertise and do not see the economic value of doing things differently
3. Defense contractors who have pockets of expertise that craft appropriate solutions to meet requirements to which they are contractually obligated
4. System integrators who are always rushed to get products to market and lack expertise to build in security and address it after deployment through patching and other means

Serge further stated that security automation will help 1) address the expertise gap of the mid-size and defense contractors and also 2) reduce overall cost and effort across all segments.

Following that, Serge provided an attack surface reference model for SoC/application-specific integrated circuits (ASICs) that examined the overall threat space for software, hardware, and software-hardware interfaces. He noted that security is difficult due to the lack of appropriate standards and a connected ecosystem, which leads to a lack of urgency and essentials. He contrasted this business problem to selling medicine: "Security is like selling vitamins — much harder than selling something like heart medication. It's largely dependent on fear (liability) versus greed (area, speed, power). Not a good space to be in." He offered a solution to infuse appropriate standards and regulations to tilt the equation and provided a technically implementable solution by means of his vision. He further elaborated by examining the digital broadcasting market ecosystem and drawing parallels to the semiconductor market space. Serge concluded by stating, "A supply chain trusted ecosystem alliance is essential for security."

3. Summary and Road Ahead

The workshop convened a diverse array of knowledgeable individuals in the field who each brought unique expertise and insights. Through collaborative discussions and presentations, these experts offered valuable perspectives and in-depth analyses on the subject matter. Their contributions not only enriched the dialogue but also provided a comprehensive understanding of the topic from various angles.

1. Representatives from semiconductor companies discussed the proactive measures being taken to bolster security and instill trust within the industry. Deliberations centered on current insights, existing challenges, and the advancements sought by stakeholders.
2. Academic scholars discussed emerging threats and their ongoing research endeavors within academic institutions. Their discourse shed light on the evolving landscape of potential vulnerabilities and efforts to address them.
3. SDOs described their efforts dedicated to formulate robust standards that elevate security, traceability, and reliability across various sectors.
4. The government underscored its commitment to foster a conducive environment that effectively mitigates risks and to enact policies that recalibrate the risk equilibrium.

In consultation with relevant experts and SMEs, NIST has identified the following next steps:

- **Security for Semiconductors**
    - Strengthen semiconductor manufacturing through the development and adoption of a ***NIST CSF 2.0 Community Profile for Semiconductor Manufacturing*** with the community (e.g., SEMI, SIA, government, academia).
    - In collaboration with the community, investigate and leverage existing standards and best practices to develop a ***Secure Life Cycle Framework for Semiconductors*** across the supply chain, including a strategy, roadmap, appropriate recommendations that focus on semiconductor supply chain traceability and provenance, and the adaptation of current software vulnerability and patch management practices for semiconductors.
- **Metrology** — Research and formulate practical ***cybersecurity measurements and metrics for semiconductors*** to continuously inform the verification and testing of countermeasures throughout the life cycle.

NIST is also investigating engagement mechanisms that leverage existing NIST and industry standards, guidelines, resources, and expertise to cultivate trust in semiconductors, such as public working sessions and a consortium to advance these initiatives in collaboration with industry and SDOs.

Appendix A. Workshop Agenda

Introduction and Overview	
9:00 – 9:25 ET	Sanjay Rekhi – NIST Kevin Stine – NIST
Hardware Development Life Cycle	
9:30 – 10:30 ET	Jonathan Ring – Office of the National Cyber Director Adam Golodner – Advisor to the Semiconductor Industry Association Matt Areno – Intel Michael Ogata – NIST
10:30 – 10:45 ET	Break
Metrology	
10:45 – 11:45 ET	Lok Yan – DARPA Mark Tehranipoor – University of Florida Jason Oberg – Cycuity, Inc. Nelson Hastings – NIST
11:45 – 12:45 ET	Lunch
Hardware/Silicon Testing	
12:45 – 13:45 ET	Mike Borza – Synopsys Niels Samwel – PQShield Ankur Srivastava – University of Maryland Rick Kuhn – NIST
Vulnerability Management	
13:45 – 14:45 ET	Dan O'Loughlin – Qualcomm Jeremy Bellay – Battelle Peter Mell – NIST A.J. Stein – NIST
14:45 – 15:00 ET	Break
Standards	
15:00 – 16:00	Jeremy Muldavin – Aerocyonics (SAE-G32) Andy Seward – TEL (SEMI) Jennifer Lynn – IBM (SEMI) Kim Schaffer – NIST
Next Steps	
16:00 – 16:45 ET	Serge Leef – Microsoft Sanjay Rekhi – NIST