IST POLICY MEMO | June 2026

Driving AI Transparency

Supply-and Demand-Based Paths Toward AIBOM

By Allan Friedman and Nick Leiserson

This policy memo calls for greater transparency in AI systems and proposes the use of Artificial Intelligence Bills of Materials (AIBOMs) to support supply chain assurance and resilience.

We present two complementary practical courses of action for policymakers that can be adopted today to encourage AI vendors to better manage risk by providing visibility into the models, datasets, software, and services incorporated into their products and services. We also propose a path toward broader industry consensus and future standardization efforts.

Introduction

Artificial intelligence systems are rapidly becoming critical components of modern society, supporting applications ranging from healthcare and transportation to cybersecurity, critical infrastructure, scientific research, and national defense. The increasing dependence on AI systems across organizations, including almost every facet of the U.S. government, makes understanding the origins, dependencies, and supply chains of those systems an essential component of trust, resilience, and risk management.

Transparency and trust are fundamental building blocks of supply chain security. Supply chain security for advanced technology is vital to fostering resilience against cyberattacks, fraud, operational disruptions, and emerging risks. Identifying components and assets is the starting point for most cybersecurity programs.[1] In the software ecosystem, this takes the form of the Software Bill of Materials (SBOM), "a formal record containing the details and supply chain relationships of the various components used in building the software."[2] Similar concerns about transparency and provenance in semiconductor supply chains have led to discussions around Hardware Bills of Materials (HBOMs). These concepts have increasingly been extended to AI systems through the idea of an Artificial Intelligence Bill of Materials (AIBOM), sometimes described as an "SBOM for AI." In the past few years, discussions around AIBOMs have accelerated across cybersecurity communities, standards bodies, government agencies, and the supply chain solutions industry.[3]

1 See, e.g. National Institute of Standards and Technology, "Framework for Improving Critical Infrastructure Cybersecurity, Version 1.0. NIST Cybersecurity Framework (CSF) 1.0," February 12, 2014, https://csrc.nist.gov/pubs/cswp/1/cybersecurity-framework-v10/final.

2 Cybersecurity and Infrastructure Security Agency (CISA) et al., "A Shared Vision of Software Bill of Materials (SBOM) for Cybersecurity (Joint Guidance)," September 2025, https://www.cisa.gov/sites/default/files/2025-09/joint-guidance-a-shared-vision-of-software-bill-of-materials-for-cybersecurity_508c.pdf.

3 See, e.g. CDX, SPDX, CISA 2023, CISA use cases, G7, UK TAIBOM

info@securityandtechnology.org · 195 41st St, #11045, Oakland, CA 94611 · securityandtechnology.org

This policy memo argues that a future in which organizations can effectively manage AI supply chain risk will require AIBOMs. Widespread adoption of AIBOMs will require progress on both the demand side—the customers consuming AI products and services—and the supply side—the creators of those products and services. On the demand side, norms, contracts, procurement requirements, and regulation can create expectations that organizations developing and deploying AI systems should understand the components incorporated into those systems. On the supply side, guiding the development of tools, processes, and organizational practices that make AI transparency achievable at scale will require a consensus vision of the minimum elements of AIBOMs and standards.

Background

Why do we want or need AI system transparency? In practice, AIBOMs could be used to solve a range of risk management challenges. At a high level, the rationale for transparency can be understood through three broad lenses: strategic, economic, and operational.

- **Strategic:** Making it clear that supply chain concerns and a risk management approach apply to the data and models that underlie the AI systems that are deployed in critical applications across the digital ecosystem.
- **Economic:** Improving information symmetry in markets, thereby creating incentives for organizations to understand and improve the provenance and quality of AI components while minimizing the use of poor quality or poorly-sourced software, data, and models.
- **Operational:** Enabling rapid assessment of and response to newly discovered vulnerabilities, integrity concerns, or other risk-relevant information.

Across all three lenses, it is clear that transparency can help to achieve specific national security and risk management outcomes.

Further, as AI supply chain risks and attacks multiply, operational resilience will only grow in importance.

Transparency in AI systems has already been identified as a necessary requirement for managing AI risk. The NIST AI Risk Management Framework (AI RMF)[4] emphasizes understanding the context, components, dependencies, data provenance, and third-party relationships that contribute to AI system risks. The AI RMF further highlights the idea of traceability and the importance of documenting "information about data, system inputs, AI system design, model development, and system outputs." This underscores the need for organizations to maintain visibility into the origins, lineage, and dependencies of AI system components, one that has been recognized by both governments and the security community.

In 2023, the U.S. Cybersecurity and Infrastructure Security Agency (CISA) noted that, "[s]ince AI is software, AI models – and their dependencies, including data – should be captured in software bills of materials."[5] Support for AI SBOMs was echoed by the G7 countries' cybersecurity agencies in a 2025 vision statement[6] and 2026 guidelines.[7] Recently, a working group meeting through the CISA SBOM community work articulated key use cases around SBOM for AI.[8] Lastly, several distinct community-led efforts within OWASP are exploring AI transparency and open-source AIBOM tooling, reflecting growing interest in AI supply chain transparency throughout the security community.[9]

4 National Institute of Standards and Technology, "Artificial Intelligence Risk Management Framework (AI RMF 1.0)," NIST AI 100-1, January 2023, https://doi.org/10.6028/NIST.AI.100-1.

5 Christine Lai and Jonathan Spring, "Software Must Be Secure by Design, and Artificial Intelligence Is No Exception," Cybersecurity and Infrastructure Security Agency, August 18, 2023, https://www.cisa.gov/news-events/news/software-must-be-secure-design-and-artificial-intelligence-no-exception.

6 G7 Cybersecurity Working Group, "A Shared Vision of Software Bill of Materials for Artificial Intelligence," May 2025, https://www.acn.gov.it/portale/documents/d/guest/paper_sbom-for-ai_19may2025_-clean-2.

7 G7 Cybersecurity Working Group, "SBOM for AI: Minimum Elements," 2026, https://www.bsi.bund.de/SharedDocs/Downloads/EN/BSI/KI/SBOM-for-AI_minimum-elements.pdf.

8 CISA SBOM for AI Tiger Team, "SBOM for AI Use Cases," community working group report, accessed June 2026, https://github.com/aibom-squad/SBOM-for-AI-Use-Cases.

9 Examples include the OWASP AIBOM Project and the OWASP GenAI Project's AIBOM initiative, two distinct community-led efforts exploring approaches to AI transparency, inventory, and open-source AIBOM tooling.

In Europe, the AI Act[10] establishes horizontal rules for AI systems sold in the EU market. Annex IV requires technical documentation describing system architecture, software components, training methodologies, training datasets, data provenance, and testing procedures. Although the Act does not require an AIBOM specifically, its provisions reflect a growing expectation that organizations understand and document the key components, dependencies, and origins of the AI systems they develop and deploy.

Collectively, these efforts demonstrate a growing consensus around the importance of greater transparency, traceability, and supply chain visibility for AI systems.

The Gap Today

Despite the growing chorus of entities calling for AI supply chain transparency, policymakers have yet to provide concrete guidance regarding the specific information organizations should maintain and exchange to support practical AI supply chain risk management.

Consider the G7 Cybersecurity Working Group's "SBOM for AI: Minimum Elements." This document makes a valuable contribution by identifying categories of information that may be relevant to AI supply chain transparency. However, despite the similar name, it deviates from the goal laid out in the original SBOM 'Minimum Elements,' published by the National Telecommunications and Information Administration (NTIA) in 2021.

The NTIA's Minimum Elements document serves to "establish the baseline technology and practices for the provisioning of SBOMs." It offers elements that "are deemed necessary to achieve the goals" of SBOM-based transparency.[11] The G7 document, on the other hand, explicitly says, "These minimum elements are not mandatory."[12] This makes it hard to argue that the G7 publication can serve as any kind of baseline or floor. A minimum elements approach comprises clearly understood, scoped, and defined terms. Indeed, many details in the G7's "SBOM for AI" document could be part of a minimum elements approach. Other details, however, such as the document's suggested "security controls" or "model input-output properties" are, at best, open-ended, and some may involve considerable ontological or standardization work before they can be implemented consistently across organizations.[13] Lastly, as we discuss further below, this document was not built on a foundation of community input and participation to reflect the needs and capabilities of a wide range of stakeholders.

Meanwhile, the maintainers of the two data formats most commonly used to exchange SBOM and related supply chain information—SPDX[14] and CycloneDX[15]—have begun developing mechanisms for representing AI-related metadata, in what they both refer to as AIBOMs. These open source, international efforts include support for documenting models, datasets, provenance, lineage, licensing, and other information relevant to AI systems. While the specific approaches differ, both efforts have made significant progress in defining encodings for representing AI-related supply chain information. Other technical discussions around AIBOM have occurred in fora like OWASP and the United Kingdom's Techworks trade association.[16] Also noteworthy are the burgeoning number of AI security and supply chain security companies that are now delivering some real form of AIBOM capabilities to their customers today. While these efforts differ in scope and implementation, they collectively demonstrate growing demand for AI supply chain transparency. AIBOM has clearly moved beyond the pure conceptual phase.

However, the practical implementation of these capabilities remains nascent, not all data types are easily

10 European Union, Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024, "Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act)," *Official Journal of the European Union* L 1689, July 12, 2024.

11 NTIA, "Minimum Elements for a Software Bill of Materials (SBOM)."

12 G7 Cybersecurity Working Group, "SBOM for AI: Minimum Elements."

13 G7 Cybersecurity Working Group, "SBOM for AI: Minimum Elements."

14 "The System Package Data Exchange," Linux Foundation, https://spdx.dev/. Information focusing on AI data encoded in SPDX can be found at https://spdx.dev/learn/areas-of-interest/ai/.

15 "CycloneDX: The International Standard for Bill of Materials (ECMA-424)," OWASP Foundation and Ecma International Technical Committee for Software & System Transparency, https://cyclonedx.org/. Information focusing on AI data encoded in CycloneDX can be found at https://cyclonedx.org/capabilities/mlbom/.

16 See: Trustable AI Bill of Materials (TAIBOM) project, last accessed June 2026, https://taibom.org/docs/. This is the product of a working group under the Techworks Deep Tech trade association.

machine processable, and, despite significant progress in developing technical mechanisms to represent AI-related metadata, there remains no broadly accepted understanding of what "minimum elements" organizations should consistently maintain and exchange to support AI supply chain risk management.[17] Many experts have advanced numerous types of data categories and data fields to address a very broad set of use cases. Many specific types of data are ill-defined, and entire standards processes might be needed to derive the needed ontologies to make some fields machine-processable.

This challenge is not unique. One key to the progress made by the SBOM movement was a shared vision developed in an open, cross-sector, multistakeholder process. With crucial buy-in from the software industry, cybersecurity experts, the open source community, and security researchers, the movement was able to drive a consensus around what defined the minimum elements of an SBOM. This, in turn, helped shape market expectations; solutions from open source and commercial tool creators; and a number of standards and regulatory regimes in the United States and around the world.[18]

The absence of such a common, widely shared vision for AIBOM presents a significant challenge. As key aspects of society grow dependent on AI, governments are taking action to mitigate risk.[19] Yet any attempt to promote AIBOM adoption through procurement, regulatory, or industry requirements without this shared vision introduces its own risks. Uncertainty around what the minimum elements are before an artifact can be deemed an AIBOM could drive underinvestment in transparency measures or overinvestment in collecting data with limited risk management value, increasing costs and slowing AI adoption.

A common understanding will drive creation of tools, automation, and repeatable processes that can reduce costs and improve consistency. Without such alignment, organizations must rely on bespoke approaches that increase complexity and inefficiency. Finally, interoperability becomes increasingly difficult when organizations receive AIBOM data from multiple sources with uneven implementation.

Consistent expectations and implementations are essential if AIBOM information is to be exchanged, aggregated, and used effectively at scale.

The gap limiting rapid adoption of AIBOM lies between high-level policy goals and low-level technical implementations. Without greater consensus at this specification layer, organizations lack clear expectations, customers struggle to request meaningful information, and tooling vendors face uncertainty regarding what capabilities to build. At the same time, the incentives and market signals needed to drive adoption remain immature. The specification layer defines a shared vision and a path to implementation. It is built explicitly around operational needs, existing capabilities, and business processes. A common vision for AIBOMs can help guide expectations across government and industry, supporting demand for transparency while enabling the development of open source and commercial tools built around a consistent set of assumptions and requirements.

This policy memo reframes these challenges in terms of supply- and demand-side approaches. On the supply side, we focus on building a shared vision, a non-trivial task requiring buy-in from a diverse group of stakeholders. On the demand front, we explore a scaffolding approach to send a clear demand signal that can lead to contractual or regulatory requirements requiring AIBOMs.

17 See: National Telecommunications and Information Administration, "The Minimum Elements for a Software Bill of Materials (SBOM)," July 2021, https://www.ntia.gov/files/ntia/publications/sbom_minimum_elements_report.pdf; Cybersecurity and Infrastructure Security Agency (CISA), "Minimum Elements for a Software Bill of Materials (SBOM)," draft for public comment, August 2025, https://www.cisa.gov/sites/default/files/2025-08/2025_CISA_SBOM_Minimum_Elements.pdf.

18 Over 24 commercial and open source solutions signed up to present at CISA's "SBOM-Solutions Showcase" in 2024. The CyBeats "Supply Chain Security Standards, Regulations and Frameworks Tracker" has 88 SBOM-related standards and regulations as of June 2026. See: Cybersecurity and Infrastructure Security Agency, "SBOM-Solutions Showcase," last accessed June 2026, https://www.cisa.gov/resources-tools/resources/sbom-solutions-showcase; CyBeats, "Supply Chain Security Standards, Regulations and Frameworks Tracker," last accessed June 2026, https://www.cybeats.com/regulations.

19 E.g., Executive Order No. 14409, "Promoting Advanced Artificial Intelligence Innovation and Security," June 2, 2026, https://www.whitehouse.gov/presidential-actions/2026/06/promoting-advanced-artificial-intelligence-innovation-and-security/.

Supply Side: Towards an AIBOM Minimum Elements

Any discussion of AI transparency and AIBOMs is not beginning with a blank page. The cybersecurity agencies of the G7 have offered a set of 48 fields that they refer to as "minimum elements," although, again, the authors take pains to point out that this set is "not mandatory." Several commercial companies offer AIBOM creation services, at least two open source projects exist that are devoted to AIBOM creation, and the authors of this policy memo are aware of more than one major contractor who claims to already be implementing AIBOM internally as a best practice—and in possible anticipation of government requirements.

While the widely-used CycloneDX and SPDX BOM formats have enumerated a number of data fields and relationships to capture supply chain risk and operations, we believe there must exist a foundation that could serve as a solid starting point. We use the term "foundational" to denote a common baseline for what widespread implementation can look like.

The core goal is to be inclusive enough to enable key use cases around risk management while being achievable enough to garner buy-in from stakeholders across the AI ecosystem, including producers and consumers of AI systems.

It is always possible to make the case for greater data inclusion. Likewise, it will always be possible to debate the costs of generating and managing data. Returning to the SBOM example, early policy work in this space followed a mantra of "crawl, then walk, then run."[20] Building out an AIBOM that can capture the data below would help to shape a consistent floor for a producer's risk management program, while helping the producer and the customer understand basic exposures.

To help orient our approach, we have selected a subset of proposed fields to include. While opinions vary on the feasibility of these data categories, our selections are frequently highlighted in discussions around AIBOM and AI supply chain security. We will note that SBOM implementation has taught us that nothing is ever as straightforward as one might expect, given the breadth and diversity of the software used in AI. For instance, there are multiple papers written to describe the challenge of dealing with uncertainty and lack of replicability for something as seemingly straightforward as software names and identifiers.[21]

Below we present what might be seen as a foundational approach to an AIBOM with the components of an AI system that support risk management. The intent is not to fully define a complete AIBOM standard but rather to illustrate the types of information that are useful, achievable, and consistently implementable today, along the lines of the SBOM Minimum Elements. Nor is this meant to be a complete version of an AIBOM Minimum Elements, but rather serve as a starting point for community discussions around progress towards that goal.

We must begin with the understanding, as articulated by CISA in 2023,[22] that AI is a subset of software, such that *any AIBOM must include or be combined with a full SBOM of that AI system.* This would include information about the software used to orchestrate the AI system, such as agent infrastructure, middleware, or tool connectivity layers.

An AIBOM should capture relevant details about the models and datasets used for training, fine-tuning, evaluation, validation, testing, retrieval, grounding, augmentation, or other model development or operational purposes. For each suggested detail, we also discuss how straightforward it may be to obtain or derive this data, potential sources of divergence, and other considerations.

---

20 National Telecommunications and Information Administration (NTIA), "Marking the Conclusion of NTIA's SBOM Process," April 9, 2022, https://www.ntia.gov/blog/marking-conclusion-ntia-s-sbom-process.

21 See, e.g., Cybersecurity and Infrastructure Security Agency (CISA), "Software Identification Ecosystem Option Analysis," October 2023, https://www.cisa.gov/sites/default/files/2023-10/Software-Identification-Ecosystem-Option-Analysis-508c.pdf; MITRE, "Data Normalization Challenges and Mitigations for Software Bill of Materials Processing," October 2024, https://www.mitre.org/sites/default/files/2024-10/PR-24-2647-Data-Normalization-Challenges-Mitigations-Software-Bill-Of-Materials-Processing.pdf; MITRE, "Considerations for Managing Challenges in Software Bill of Materials," April 2026, https://www.mitre.org/sites/default/files/2026-04/PR-26-0685-Considerations-for-Managing-Challenges-in-Software-Bill-of-Materials.pdf.

22 Christine Lai and Jonathan Spring, "Software Must Be Secure by Design, and Artificial Intelligence Is No Exception."

Dataset Details

Note: Bolded text describes the proposed element. Italics provide author commentary on that element.

Identity

DATASET NAME

Foundational, but not always completely straightforward if a dataset has multiple names, nicknames, or truncations. (In the SBOM context, consider, for example, challenges with software identity.)

DATASET VERSION AND/OR DATE

Explicit versioning may be more difficult for highly dynamic or evolving datasets, so a time-stamped release identifier may be needed.

DATASET LOCATION - A download URI for a public dataset, or an indication that it is from a private source.

Integrity and Verification

INTEGRITY REFERENCE

Hashes can be difficult for very large datasets or collections of data, such as a collection of documents in a shared repository. As such, it may be appropriate to use a cryptographic identifier, such as a checksum, digest, or equivalent integrity mechanism, sufficient to uniquely identify the referenced version of a dataset and to indicate that it has not been tampered with, nor is identical to a dataset that has been compromised.

Governance and Usage

SENSITIVITY OF THE DATA - Is the data freely accessible, proprietary, classified, export controlled, commercially sensitive, PII, etc.?

No universal schema exists today, but could probably be defined as an extensible tag list.

LICENSE

This can borrow heavily from Open Source Software Licensing, but other approaches may exist in the data space around uses, data protections, and proprietary sources.

Provenance and Context

DATA SUPPLIER - Where did the AI system supplier directly obtain the data from?

This can be automation friendly by tracking inclusion.

DATA CREATOR - Is there a single entity that is deemed responsible for creating the original database?

Supporting data documentation - A set of resources, such as datacards or corporate catalogs, that can provide more unstructured insights into data details.

Some types of provenance and context information are frequently cited as important for AI supply chain analysis, but are considerably more difficult to capture and exchange consistently. Often, the challenge is not that the information is unavailable, but that it is nuanced, context-dependent, or resistant to a single authoritative description. Many of these fields may ultimately prove important for risk management and national security purposes, but further work is needed before they can be broadly standardized and automated.

Further provenance

DATA ORIGIN - What are the underlying sources from which the data was collected, generated, or derived?

This information may be unavailable to downstream users, difficult to verify, or challenging to represent consistently in a machine-processable fashion.

DATA LINEAGE - What existing dataset or datasets was this dataset derived from?

Understanding lineage can support provenance and risk analysis, but introduces the same challenges around identification, versioning, and integrity as discussed above.

COUNTRY OF ORIGIN

In some contexts, organizations may wish to understand the geographic origins of data. However, modern datasets are frequently assembled from multinational sources, making country attribution difficult to determine consistently and precisely.

DATA PROCESSING HISTORY - Information covering cleaning, filtering, augmentation, and curation.

These are not well defined in 2026, and if the data was processed prior to the current user's period of control, they may not have any insight into what happened.

Model Details

Note: Bolded text describes the proposed element. Italics provide author commentary on that element.

To build out an AIBOM, it is important to understand the model or models used in an AI system. Again, many of the same challenges exist around naming and identification, as well as capturing dynamic models or systems with active feedback loops. The below is a starting point.

Identity

MODEL NAME

Names provide a useful human-readable reference for a model. However, names are not guaranteed to be unique and may be reused, abbreviated, or referenced inconsistently across organizations. As such, a model name should not be relied upon as the sole means of identification.

MODEL VERSION OR RELEASE IDENTIFIER

See above for discussions about versioning. Many models follow conventional versioning practices, but not all do.

MODEL IDENTIFIERS

Models may have multiple identifiers, including registry IDs, URIs, or other unique references.

Integrity

INTEGRITY REFERENCE - A cryptographic identifier or equivalent mechanism sufficient to uniquely identify the referenced model version and support integrity verification.

See above for discussion of different approaches to maintain integrity.

Governance

MODEL LICENSE - Licensing information governing the use, modification, distribution, or deployment of the model.

Open source licenses are increasingly common, though commercial and proprietary licensing arrangements may also apply.

Provenance and Context

MODEL SUPPLIER - The organization or service from which the model was obtained.

MODEL ORIGIN - The organization that originally created, trained, or published the model.

This field helps distinguish the original source of a model from subsequent distributors, hosts, or integrators.

MODEL LINEAGE - A model or set of models from which this model was derived.

Understanding lineage is critical for tracing newly identified risks or inherited characteristics, but may be difficult to document precisely in cases involving extensive retraining, model merging, or other significant modifications.

SUPPORTING MODEL DOCUMENTATION - References to model cards, evaluation reports, more complete safety analysis, or other relevant information, although it may not be directly machine readable.

Other fields may provide useful context for understanding a model's role within an AI system, but are less mature, inconsistently defined, and tied to rapidly evolving technologies and implementation patterns. As a result, these fields may be valuable for future AIBOM efforts but are less suitable for broad requirements today.

Further context

MODEL ROLE - What function is the model serving? This could include foundation models, fine tuned models, adapters, ensemble models or embedding models.

A finite, extensible list seems possible.

Not all the data above may be available to an AI system producer. In some cases, it may be unknown. In other cases, it may be because the supplier chooses not to disclose it. This is particularly true for provenance data. Many of the largest commercial foundation models treat training data and lineage as trade secrets, which means that the data origin, lineage, and country of origin fields above, often relevant to security, will frequently be unavailable from precisely the suppliers whose models are most widely used. However, this is not a reason to

leave this field off of a Minimum Elements specification. Instead, the absence should be documented explicitly as a known unknown. This could be an example where inclusion in a widely-used approach could help shift norms towards greater transparency, especially when combined with demand-side pressure.

In addition to the above list, many other types of metadata around AI systems have been proposed. Some will require more discussion to build out common vocabularies or ontologies, such as the intended uses for a model or the uses that have been deemed inappropriate or out of scope for a particular model. An enumeration of security controls and guardrails would similarly require further work to determine how to bound this list and effectively describe controls so that they can be understood by customers. Others are more holistic properties of an entire AI system which may be important but fall outside the primary scope of supply chain components.

One important limitation of this approach is that it focuses primarily on the components, dependencies, and provenance of AI systems rather than their real-time behavior. Modern AI systems are increasingly dynamic, incorporating external data sources, retrieval mechanisms, tool use, environmental inputs, and feedback from users and other systems. Emerging agentic architectures and boundary-blurring protocols like Model Context Protocol (MCP) further complicate this picture. Recent efforts, such as the Coalition for Secure AI's work on agentic identity and access management and MCP security, illustrate how agents may be instantiated dynamically, operate with delegated authority, interact with external tools and services, and then cease to exist when their task is complete.

As a result, an AIBOM can provide important visibility into what an AI system is built from, but it cannot fully capture the state or behavior of that system at a particular moment in time. AIBOMs are important, but they should not be construed to solve—or surface—all AI security concerns.

The broader AI security community is actively developing approaches to address assurance, monitoring, and governance of these dynamic systems, including techniques for managing supply chain risks that emerge during operation rather than during development or deployment. While these efforts are promising, they remain relatively nascent and are not yet suitable as a universal requirement to be imposed on the diverse range of AI systems used today.

Lastly, this foundational approach should not be seen as the limit of what data can or should be captured. We fully expect AIBOM fields to grow over time, driven by a better understanding of what data is needed to understand the risks of existing and emerging AI systems, who can use it, and how that data can be generated and encoded effectively.

Demand Side: Maintaining AI Supply Chain Data

Rather than focusing on developing a shared vision of AIBOM directly, an alternative approach is to look at the demand side of the equation. If government regulations and industry procurement norms begin requiring AIBOM, that requirement would then drive market players to develop transparency solutions. However, without AIBOM minimum requirements, defining what specifically is required in a contract or regulation could remain challenging. Beyond a specific AIBOM requirement, high-level visions of AI security, as articulated above, emphasize the importance of tracking supply chain data. What is the path that gets us to a world with greater supply chain risk awareness and better practices? More to the point, how can policymakers and downstream risk owners drive demand for greater transparency? The good news is that we can start this today, without consensus minimum requirements, in a way that will still encourage that consensus to form.

The demand side begins with some form of forcing function or requirement that organizations understand what is in the products they manufacture and sell.

Internal visibility is a basic component of supply chain risk management. Turning this basic concept into a requirement in government procurement language, business-to-business contracts, or sector-specific regulation codifies this basic hygiene practice while avoiding being overly prescriptive, especially given the current lack of consensus around minimum requirements for a shareable BOM. At its core, a requirement would simply mandate that a producer must know what is in its products, be able to answer questions about potential risks, and understand specific uncertainties in its supply chain. Because such a mandate centers around outcomes,

rather than requiring any specific mechanism or standard, it represents a more flexible starting point for AI supply chain security.

What would such a requirement look like? It starts with a mandate to track AI system components. Other policies in the supply chain space contain language that require tracking system components without any formal structure or sharing requirement. For example, the payment card industry has produced a non-governmental standard governing the security of systems that handle credit card data. The Payment Card Industry Data Security Standard 4.0 requires "an inventory of bespoke and custom software, and third-party software components incorporated into bespoke and custom software."[23] This requirement is lightweight, in that it does not go into details about the exact nature of how components should be tracked, or what precise details should be kept. In the AI context, such a requirement could easily be expanded beyond software to encompass system components, including data and models.

Of course, the purpose of maintaining AI supply chain information is not compliance for the sake of compliance. First, maintaining AIBOM can augment an organization's internal risk management processes. Simply understanding what components are incorporated into an AI system, where those components originated, and how they relate to one another can improve governance, support procurement and sourcing decisions, identify potential concentrations of risk, and provide a foundation for future security and assurance activities.

Once an organization has some understanding of their AI supply chain, their customers and regulators may actually ask them to do something with this data. We identify two general types of requests that policymakers and customers might make. First, they may ask for the organization to respond to novel risk information. Second, they may ask the organization to respond to specific queries about the presence of any component or type of component.

Turning first to integrating and responding to risk information, there are a host of data sources that could come into play, ranging from known vulnerabilities to identified threats and information about the trustworthiness of organizations or even specific individuals. One of the most insidious software supply chain attacks in recent years was traced to the actions of a single (still anonymous) open source contributor.[24] Such a requirement could ask organizations to track risk data and respond appropriately. This includes understanding whether they are affected, how they are affected, and what the optimal mitigation should look like.

The second type of requirement could ask organizations to respond to queries or requests about specific supply chain components, actors, or sources. Again, we see this in other supply chain policies. Congress established a requirement in 2023 that all government contractors must certify that they do not use hardware components from three specific Chinese companies.[25] In the AI context, a company in Europe, operating under strict privacy laws, may be averse to using AI systems built on datasets collected outside the context of a data protection regime. Rather than demanding full details about the data source, the company may ask for attestations, certifications, or even third party assessments to ensure that the data used to build these systems meets specific requirements. In this case, demand for tracking supply chain data would be driven by the need for the capability to respond to this kind of request or directive. AI system vendors that maintain high quality metadata in a machine-processable form will find it easier and cheaper to respond to this demand.

What data would an organization need to track internally to meet these use cases? Such tracking must go beyond a standard inventory with basic information to identify each component. For more rigorous requirements, an organization should have a core understanding of provenance, including development history and potential exposure to foreign ownership, control, and influence.

Transparency includes documenting uncertainty. For complete implementation, an organization must also clearly track when they do not have enough information to make relevant risk-based decisions about the supply chain. This includes software of unknown provenance, data sets or data sources with insufficient origin information, or models that have been modified by unknown processes or by unknown actors. In other words, "known

23 PCI Security Standards Council, "Payment Card Industry Data Security Standard: Requirements and Testing Procedures," Version 4.0, March 2022, https://www.pcisecuritystandards.org/document_library.

24 Piotr Przymus and Thomas Durieux, "Wolves in the Repository: A Software Engineering Analysis of the XZ Utils Supply Chain Attack," in Proceedings of the 22nd IEEE/ACM International Conference on Mining Software Repositories (MSR 2025) (2025) https://doi.org/10.1109/MSR66628.2025.00026

25 U.S. Congress, National Defense Authorization Act for Fiscal Year 2023, Pub. L. No. 117-263, § 5949, "Prohibition on Certain Semiconductor Products and Services," 136 Stat. 2395, 3458–61, 2022.

unknowns" must be explicitly captured.

There are risks with taking a demand-centric approach that stops short of requiring an actual AIBOM. Since the data is not shared with customers, only kept internally, the scheme would rely on attestations that such data exists and is sufficient for the purposes described. Existing authorities for false claims and contract law can help mitigate this risk if the requirements are clear enough. More dynamic third party certifications can also be used, even without a formal standard or conformance regime.

To address the operational mandate described above, this approach would also require clear communication about how an organization must respond to novel risks, what types of risks and threats intelligence would be covered, and where an organization can find this information. None of these details are impossible to define, although both the data and the response should be articulated in a manner that can be updated as needs and capabilities evolve.

By creating an obligation to maintain AI supply chain information, this approach creates demand for processes and practices around supply chain data. Organizations may satisfy these requirements using a range of data, including SBOMs, internal repositories, development pipelines, enterprise resource planning (ERP) systems, or future approaches. Meeting these requirements effectively and efficiently would require automation. This, in turn, requires tooling. For the tooling ecosystem to scale more broadly the bespoke internal tools, we will need more commonly implemented mechanisms of capturing and using AI supply chain data. In the short term, demand can also drive attention to AIBOM efforts overall. Over time, this demand can help drive convergence toward shared AIBOM practices and standards developed through the multistakeholder process described below—and eventually allow for AIBOM to be required directly.

Consensus before Regulation: A Roadmap to AIBOM

How do we solve for the *specification* layer? Is it by driving demand or increasing supply? While either strategy could work, there are risks to adopting either in isolation.

A supply-side approach risks petering out or failing to draw in all of the requisite stakeholders because it could be viewed as an "academic exercise." Without a clear source of demand for AIBOM, the supply-side approach also risks becoming a lengthy process. Building consensus without a deadline is always challenging. Doing so in the face of the incredibly rapid rise in AI capabilities—and corresponding integration of AI into core aspects of society—could be lethal.

However, a demand-side approach also poses significant risks. Setting a requirement that AIBOMs be included as part of a procurement (or other regulation) by a date certain without a consensus as to the minimum elements of an AIBOM risks fragmenting the landscape. Dealing with non-standardized AIBOMs increases the burden on vendors, who have to attempt to map different requirements and then present similar data to different customers. This approach also makes it more difficult for the customers themselves, who need bespoke systems and training to handle non-standard AIBOMs. Both of these factors increase the likelihood that industry will work to oppose immature requirements altogether, rather than organizing to develop consensus.

On the other hand, a demand-side approach that only requires internal supply chain monitoring, without specifying an AIBOM, will likely fail to solve policymakers' long-term needs. The goal is putting information directly in users' hands, rather than requiring independent vendor queries as threats emerge. What's more, relying on self-attestation based on intentionally vague requirements risks failing to adequately incentivize data collection and retention.

De-risking involves a comprehensive process that links the supply and demand sides. On the supply front, this means encouraging the creation of a multistakeholder working group to translate policymaker priorities and needs into specific minimum elements of an AIBOM, implementable in common data formats. On the demand side, it means setting a clear, phased timeline for full AIBOM adoption. Interlinking these two workstreams creates a positive feedback loop: the lessons from early data collection directly inform the creation of the AIBOM specification, which in turn builds confidence in more stringent requirements. Policymakers can take key steps to support this linked strategy.

A Multistakeholder Approach

The complexities of AIBOM development demand broad, cross-sector engagement. AI systems are increasingly used across government and critical infrastructure, including the healthcare, finance, transportation, and manufacturing sectors. Many of the same models, datasets, tools, and services are shared across these communities, creating common supply chain risks and common opportunities for transparency. A successful AIBOM effort should therefore engage model developers, AI system integrators, cloud providers, software suppliers, cybersecurity practitioners, standards bodies, researchers, government agencies, and end users. Broad participation will help ensure that any resulting specification, and corresponding model contract language, reflects practical implementation realities and remains relevant across procurement or regulatory contexts.

A trusted and neutral nonprofit organization could play an important convening role in this process. Such a forum could provide a venue for developing consensus around AIBOM concepts, identifying implementation challenges, coordinating with existing standards efforts, and establishing a roadmap for future evolution as AI technologies and supply chain risks continue to mature. In keeping with the analytical framework, the multistakeholder process should encompass working groups covering both the supply and demand sides.

Legislation or administrative policy can support this approach in several ways. Endorsing the need for a multistakeholder process demonstrates potential demand for its output, as does funding entities that can pull together the relevant parties. Policymakers can also direct government agencies to participate and encourage industry representatives to contribute by recognizing their service to the broader community.

SPECIFICATION WORKING GROUP

To support the supply side of the process, policymakers should continue to refine their objectives at the strategy layer: what problems are they looking for AIBOM to solve? In addition, they should set clear red lines to define the bounds of the process, articulating where policy has already been specified to a degree that additional input is not required outside of traditional legislative or administrative processes. To that end, we suggest some elements of a specification working group charter.

1. **AIBOM should build on existing BOM approaches.** SBOMs are a non-negotiable part of supply chain data, and should therefore be part of any requirement. Hardware components also play a key role in the supply chain of AI systems. While hardware provenance discussions are less mature than software or AI, HBOM should be an explicit part of any AI supply chain security roadmap.[26]
2. **AIBOM data should be readily available to the AI system supplier or contractor.** Suppliers should be reluctant to use components for which the very basic identification, provenance, and integrity data is not readily obtainable. The specification should reflect an awareness of how a supplier might obtain this data, ideally in an automated and verifiable fashion.
3. **AIBOM data should also have direct utility to the mission of supply chain assurance and resilience.** While there is a host of metadata one may wish to know about building and using an AI system, this approach to an AIBOM is specifically oriented around cybersecurity threats and responses. An AIBOM must have enough detail to detect and respond to novel risks, threats, and incidents.
4. **An AIBOM specification should aspire for consistency of implementation.** Organizations will create AIBOMs for multiple systems, often across divisions inside the organization. Customers will receive AIBOMs from many suppliers. In both cases, the AIBOMs should look similar enough to enable the use of automation for detecting and managing risks. Tooling vendors require sufficiently consistent data to build out practical intelligence and actions based on AIBOM data. The specification should map cleanly to both of the widely-used BOM formats.

Any effort to build out an AIBOM specification should not seek to replace or compete with existing standards efforts, such as SPDX and CycloneDX. Instead, it should identify a common set of data elements, relationships, and expectations that can be implemented through those existing frameworks and their ongoing AI-related work, such as those discussed in Supply Side: Towards an AIBOM Minimum Elements. As with the original SBOM Minimum Elements effort, the objective is not to mandate a particular format, but to establish a common understanding of what information is important and how it can be used to support supply chain assurance and resilience.

26 For more information on HBOM, see HBOM.tech.

ADOPTION WORKING GROUP

Policymakers also have opportunities to support an adoption working group. The adoption group should focus on mapping procurement and regulatory opportunities to improve AI transparency. A key feature of BOMs, whether for software, cryptography, hardware, etc., is that it is relatively easy to determine if a BOM exists and is well-formed. This reduces costs of compliance. It also improves the efficacy of requirements, whether from buyers or regulators, as a deviation from the BOM discovered in a product or service can represent a breach of contract on its face. Interfacing with purchasing and legal teams would give the working group additional insight into how to craft model contract language for requesting AIBOMs that meet emerging minimum requirements. Models like this have been effective for a range of issues, including ensuring student privacy in education technology,[27] driving use of cybersecurity education frameworks,[28] and developing approaches to cyber incident reporting.[29]

The adoption group should also encompass workstreams related to the use of AIBOMs. There is reason to believe that, as with SBOMs, the uses will not be static. The goals of SBOM have been laid out over time, and new use cases have emerged as the processes have matured and more organizations explore what to do with data. The new structured datasets represented by AIBOMs do not, in and of themselves, significantly improve customer supply chain security. They do, however, enable understanding and rapid response to new information about supply chain risks. Iterating with supply chain security teams as minimum requirements are solidified will help ensure that there is a nascent AIBOM doctrine ready to take advantage of these new data.

Crucial to the success of the working group will be its ability to drive consensus on a maturity model and associated roadmap for AI supply chain transparency. A concrete timeline gives space for norms to develop and vendors to adapt while also allowing organizations to build workflows related to collecting or ingesting AI provenance data.

Beyond actively participating in the working group, policymakers can best support the effort by committing to sustained, phased improvements in AI supply chain transparency, culminating in an AIBOM requirement. Policymakers can also coordinate across international borders to drive regulatory alignment from the outset, rather than attempting to harmonize disparate approaches after they have already been implemented.

Conclusion and Summary of Recommendations

The need for improved supply chain risk management around AI systems is clear. AIBOM represents a promising conceptual framework that, by building on SBOM concepts, can provide more transparency about AI systems to users. In turn, users can proactively manage risk and respond quickly to emerging threats and vulnerabilities.

However, while technical formats exist that can convey AI system metadata, there remains a gap in the specification layer, as there is not yet consensus on the minimum elements of an AIBOM. Both supply-side and demand-side interventions could drive adoption of a minimum elements approach. However, because of the interplay between them, the best outcomes will follow from a coordinated, multistakeholder process that encompasses both workstreams. Policymakers should look to incentivize the creation of such a process and participation from both government and non-government stakeholders.

For policymakers in particular, we offer the following recommendations:

Principles for AIBOM

In any work on AIBOM, policymakers should consider these principles:

» **The AIBOM specification should explicitly address the availability, utility, and consistency of the AIBOM data**. Weighing all three criteria is critical for building a practical, scalable, effective AIBOM.

27 Student Data Privacy Consortium, "National Data Privacy Agreement," April 24, 2024, https://privacy.a4l.org/national-dpa/.

28 National Institute of Standards and Technology, "CyberSeek," created January 10, 2017, updated March 25, 2025, https://www.nist.gov/itl/applied-cybersecurity/nice/resources/cyberseek.

29 Institute for Security and Technology, Cyber Threat Alliance, et. al., "Cyber Incident Reporting Framework," November 2022, https://securityandtechnology.org/wp-content/uploads/2024/10/Cyber-Incident-Reporting-Framework-CTA_IST.pdf.

- **Any AIBOM must include an SBOM**. SBOMs are widely accepted and rapidly maturing. Discussions of AIBOMs should start with the SBOM as a foundation and build atop it.
- **HBOMs should be integrated into the approach**. While out of scope for this memo, hardware supply chain risks are increasingly relevant as the global AI race accelerates. AIBOM development should contemplate and complement HBOM efforts.
- **AIBOM will continue to evolve.** As AI technology continues to develop, the way system components should be captured will evolve as well. Language should avoid locking in specific versions of an AIBOM Minimum Elements and any processes intended to drive consensus should be planned with extensibility in mind.

A Supply and Demand Strategy

To accelerate the development of AIBOM, policymakers should consider focusing on:

- **Adhering to a supply and demand strategy**. Focusing on only one side of the equation poses significant risks to achieving unfragmented AIBOM adoption in the near term.
- **Embracing a multistakeholder process led by an independent, civil society organization.** A successful AIBOM effort should engage model developers, AI system integrators, cloud providers, software suppliers, cybersecurity practitioners, standards bodies, researchers, government agencies, and end users.
- **Encouraging participation in AIBOM development processes.** Policymakers can direct government agencies to participate and encourage industry representatives to contribute by recognizing their service to the broader community.

Supply Side Activities

A significant gap remains in the specification layer. Policymakers looking to close that gap should consider:

- **Setting a goal of having an AIBOM Minimum Requirements document**. Modelled after work with SBOM, this could sit atop multiple data standards. *Sample language relevant to this recommendation can be found in the Appendix.*
- **Continuing to refine desired outcomes at the *strategy layer*.** These will be key inputs for a multistakeholder process and should continue to build on work at the G7.

Demand Side Activities

Policymakers have unique authorities and capacity to influence the demand side of AI supply chain transparency. In doing so, they should consider:

- **Committing to sustained, stepwise, improvements in AI supply chain transparency, culminating in an AIBOM requirement**. A phased approach creates a positive feedback loop with the demand side. *Sample language relevant to this recommendation can be found in the Appendix.*
- **Coordinating across international borders to drive regulatory alignment from the outset.** This avoids having to attempt to harmonize disparate approaches after they have already been implemented.

About the Authors

Dr. Allan Friedman is Senior Adjunct Technical Advisor at the Institute for Security and Technology and a cybersecurity policy expert whose work has helped shape modern approaches to software supply chain transparency. During more than a decade in federal service at NTIA and CISA, he led efforts to advance Software Bills of Materials (SBOMs) and related cybersecurity initiatives. His current work focuses on transparency across software, hardware, and AI supply chains, and he advises industry on related risk management practices as Technologist-in-Residence at TPO.group.

Nicholas Leiserson is the Senior Vice President for Policy at the Institute for Security and Technology (IST). A legislative strategist and technologist, he has spent 15 years addressing cybersecurity risk and resilience and managing multidisciplinary teams of senior professionals at the White House and on Capitol Hill.

Appendix - Model legislative language

Potential supply-side language

1. *Within 540 days, the [appropriate Official] shall finalize a document outlining the minimum elements for an AIBOM for AI systems acquired by the [relevant Departments or Agencies] .*
2. *The document outlined in a) shall cover elements about the models and datasets used for training, fine-tuning, evaluation, validation, testing, retrieval, grounding, augmentation, or other model development or operational purposes.*
3. *In developing the document outlined in a) the [Official] shall consider data fields including:*
    a. *Dataset name*
    b. *Dataset version and/or date*
    c. *Dataset location*
    d. *Integrity reference*
    e. *Sensitivity of the data*
    f. *License*
    g. *Data supplier*
    h. *Data creator*
    i. *Data origin*
    j. *Data lineage*
    k. *Country of origin*
    l. *Data processing history*
    m. *Model name*
    n. *Model identifiers*
    o. *Model version or release identifier*
    p. *Model supplier*
    q. *Model origin*
    r. *Model lineage*
    s. *Model license*
    t. *Integrity reference*
    u. *Supporting model documentation*
4. *AIBOMs specified in the document outlined in a) must be machine-processable and human-readable.*
5. *The [Official] shall convene a multistakeholder working group, including representatives from academia, civil society, and diverse industries, including [specific sectors], to develop the document outlined in a).*

Again, we do not recommend locking in specific data fields in statute, particularly for such dynamic technology. A requirement to "consider" these fields helps both the appropriate official and the broader AI community have a common starting point for the multistakeholder effort. To make sure that this effort feeds directly into new contracts, one might consider language such as:

6. *"Within 180 days of the creation of the document outlined in (a), the [appropriate Official], shall require the [appropriate regulation] to be updated to require that all [contracts, approvals, or other regulatory mechanism involving an AI system] contracts involving the procurement of an AI system include an AIBOM consistent with said document."*

Potential demand-side language

1. *Contractor shall maintain an inventory and associated records, such as a well-structured, machine-readable AIBOM, of all components incorporated into or configured as part of AI systems provided under a given contract, including software, models, datasets, retrieval resources, tools, and services.*

2. *This inventory must be sufficient to enable a timely assessment of the impact of newly identified vulnerabilities, security risks, integrity concerns affecting software, models, or data, and other newly available risk-relevant information affecting components incorporated into or relied upon by the AI system.*
3. *Contractor shall maintain a machine-processable SBOM that captures all software components and supports automation and vulnerability management.*