Jointly drafted by BSI and ACN

# A SHARED G7 VISION ON SOFTWARE BILL OF MATERIALS FOR AI

## Transparency and Cybersecurity along the AI Supply Chain

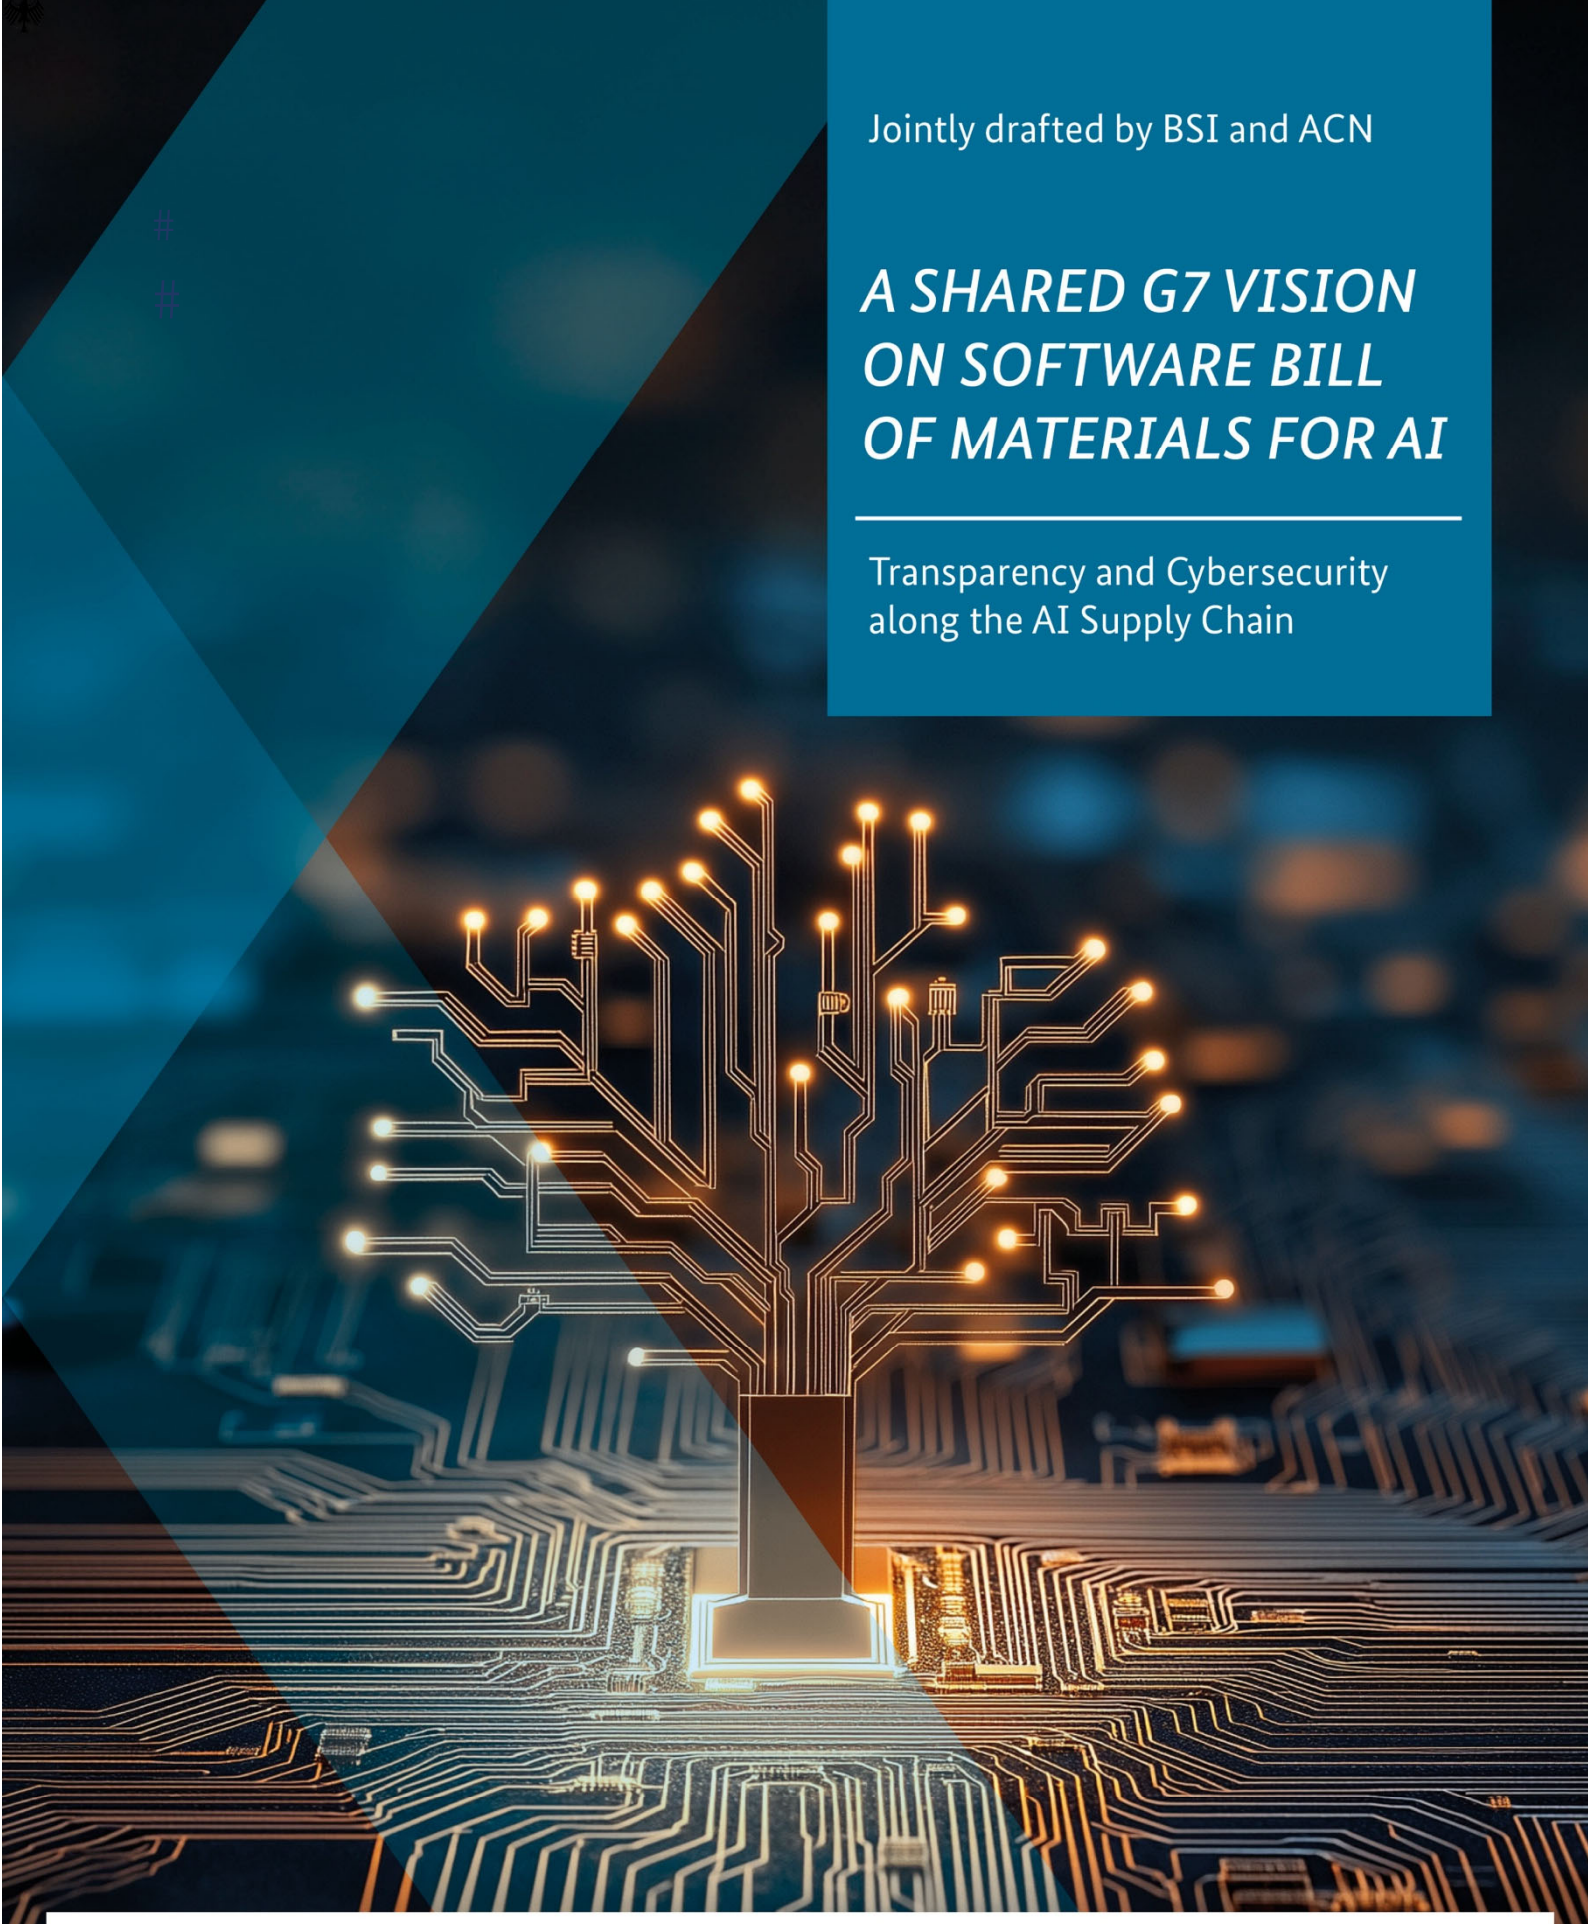

Federal Office
for Information Security

G7 2025
KANANASKIS

AGENZIA PER LA CYBERSICUREZZA NAZIONALE
ACN
SAPIENTIA ADIUVAT

#

*This paper has been written under the work stream "Smarter Together: Artificial Intelligence" of the G7 Cybersecurity Working Group. It serves as food for thought and does not intend to contradict work conducted in existing G7 groups such as the Hiroshima AI Process. In some jurisdictions a number of elements of SBOMs for AI proposed in this paper may be expected to be covered, for example, through legal requirements and obligations or existing or forthcoming standards. Due to the ongoing evolution of legal and policy frameworks within the G7 members, this paper is open for further evolution.*

# 1. Introduction

Organizations, institutions and the general public all around the world are using AI systems for multiple purposes. Developing such systems is highly complex: it requires a vast amount of resources (e.g., energy, data), infrastructure (esp. GPUs), and human expertise. Many AI systems often rely on existing base models - either commercial or open source - and are adapted according to their application purpose.

Like most modern software systems, AI systems are complex. Complexity often leads to a lack of insights into how exactly an AI system works and which components and elements an AI system is based on.

As a result, key cybersecurity issues, such as potential weak points, vulnerabilities, manipulations or compromises, are difficult to detect. Indeed, aspects such as how an AI system has been trained, including the data used and the underlying base model, are key to ensure trustworthiness, security and safety of AI systems, whereas we note that safety and security of AI systems are related as described for example in G7 Leaders' Statement on the Hiroshima AI Process.

The *G7 Cybersecurity Working Group* suggests introducing a common concept for Software Bill of Materials for AI (SBOM for AI).[1] An SBOM for AI consists of a structured record of details and supply chain relationships for the various components used in building an AI system. The goal of an SBOM for AI is to contribute to fostering security of AI systems through AI supply chain transparency and traceability of its components and dependencies.

This paper presents a shared vision and a first high-level summary of the SBOM for AI concept, including its benefits for cybersecurity, its properties, and an initial proposal for its example minimum elements. While SBOMs for AI are not explicitly a cybersecurity tool, if designed and used correctly in conjunction with appropriate tools (e.g., vulnerability management software), they could foster the transparent inventory needed to help secure the supply chain. The paper concludes with an outlook on necessary next steps to further proceed with the technical implementation of an SBOM for AI framework.

#

[1] The term SBOM for AI follows the already established concept of Software Bill of Material (SBOM) in the field of traditional software management.

# 2. Improving cybersecurity through transparency along the AI supply chain

One of the key **challenges of securing an AI system** is the vulnerability of its supply chain to both traditional and novel attack vectors. The depth and complexity of AI supply chains, coupled with the evolving and dynamic AI lifecycle, represents a considerable attack surface. Successful cyber attacks often aim at compromising a product before it reaches the consumer/end-user, including also AI components. The goal of so-called supply chain attacks is often to gather/steal sensitive information, pre-positioning and more generally to cause damage, either in the relationships between the stakeholders and/or financially, e.g., by tampering or poisoning data, causing unproductivity, misinformation or pursuing own interests.

Against this backdrop, **improved cybersecurity along the AI supply chain can be achieved by increasing transparency,** specifically with regard to accessing information on the creation process of the final AI system, as well as its individual components and dependencies. While for traditional software products the Software Bill of Materials (SBOM) concept may support mitigation to the above mentioned cybersecurity attacks, there is no internationally established and practically applicable common practice for systems using AI yet. In conjunction with other security tools, an SBOM for AI can increase transparency along the supply chain and thus contribute to cybersecurity.

**An SBOM for AI may bring transparency and knowledge about AI system composition**. It fosters **vulnerability management** and patching by minimizing the response time required to check if known vulnerabilities are deployed within the AI system components, supporting **risk management.**

Moreover, an SBOM for AI allows both **AI model tracking,** addressing issues related to performance while speeding up the entire **security compliance verification** process, simplifying auditing and keeping track of already existing compliance attestations. Furthermore, the usage of already proven components can reduce costs during the development phase of an AI system. As part of existing Secure by Design paradigms, a persistent use of an SBOM for AI can have positive effects on expensive and time consuming reworkings, such as model retraining, or damage repair. An SBOM for AI also facilitates license management. Most importantly, an SBOM for AI strengthens the autonomy and awareness of AI systems stakeholders by enabling them to make a carefully considered decision whether an AI system, an individual component or indeed a supplier is suitable for a particular purpose or not.

#

# 3. Software Bill of Materials for AI

## *Properties*

To allow an SBOM for AI to be effective, it needs to ensure that the **three following properties are satisfied**:

- being able to capture the **static and dynamic aspects of AI systems** (e.g., datasets used for training, testing and validation during the lifecycle of the system or learning outcomes) that distinguish them from traditional software systems;
- being able to be easily processed **automatically** and tool generated in a **machine-readable format**;
- being able to **leverage structured data formats as much as possible**, to ensure that the relevant information is available transparently upon demand to all the stakeholders.

Furthermore, it is equally important to clearly define the **information set** that an SBOM for AI should include, defined as its **"minimum elements"**.

## *Example minimum elements*

An SBOM for AI should be composed of a set of minimum elements to capture the distinctive features of an AI system, ensuring compatibility and providing an adequate level of transparency for all the stakeholders. It should automatically build upon information captured by each of the AI components, providing an understanding of the flow between the AI elements of the system. While some transparency mechanisms exist, **this effort aims to highlight a core set of data fields that are machine-generatable and machine-processable**. It is important to highlight that these minimum elements represent recommendations to a reasonable extent and should be decided accordingly to the specific context of use. Here below is an exemplary set of **high-level minimum elements for a G7 SBOM for AI framework**, which may extend the information used for traditional software bill of materials (e.g., supplier name, component version)[2], listed as clusters that can embed more detailed information on:

- ***Models*** used by the AI system, including basic information to identify the model, describe how the model was created, and spell out how the model is intended to be used.
- ***Learning***, including the description of the training techniques and pipelines and information about training datasets in, e.g., datasheets for datasets.
- ***Datasets*** used during the whole lifecycle of the model, including basic information that documents the identity, creation, use, and provenance of data.

#

[2] As an example we can consider the information included in the US Department of Commerce NTIA document found in reference.

- ***Safety and security characteristics***, such as a link or reference to the safeguards or guardrail implementations, safety alignment, compliance attestations and cybersecurity best practices adopted during the AI lifecycle.
- ***System level characteristics***, such as a link or reference to a description of the flow between the AI elements and how the model consumes input data.
- ***Key Performance Indicators*** of an AI system, including model benchmark evaluation results.
- ***Licensing*** information about the components of an AI system.
- ***Infrastructure*** used by the AI system, including the software components specifically required to deliver an AI system.

The list is open for further expansion of the clusters in the future to keep pace with the rapid development of technology.

To increase trustworthiness and to avoid giving a false sense of security, an SBOM for AI should be verifiable as a whole. This implies not only the verification of its individual components - e.g., via cryptographic hashes or digital signatures from the corresponding manufacturers - but also of the entire SBOM for AI. In order to achieve this goal, a viable SBOM for AI should at least be digitally signed by its manufacturer. While individual components are signed within the SBOM for AI, the signature of the entire SBOM for AI has to be verifiable from the outside.

# 4. The way forward

## *Challenges*

An SBOM for AI should include the unique features that distinguish an AI system, in addition to traditional software components. Before introducing an SBOM for AI, tools like system cards and model cards have been proposed both by private companies and by AI regulation as tools to offer increased transparency into AI models. Despite their effectiveness in some contexts, today these tools suffer from lack of harmonized and machine-readable formats, automation and interoperability with other tools. Data management is also an important consideration. SBOMs are not assumed to be publicly available today by regulation or market expectation, and intellectual property protection assumptions should hold for an SBOM for AI. **Capturing the dynamic features of the AI model** through adequate file formats and fields represents a challenge for building an SBOM for AI. At the same time, an SBOM for AI needs to be able to provide traceability of training pipelines and datasets, especially in the case of proprietary closed models, synthetic data and pre-training information, where a very large number of diverse data corpuses and sophisticated data processing pipelines are used for creating base models. Furthermore, it is critical to **keep an SBOM for AI current with the speed at which AI technology develops**, adding new information and relevant fields when needed, such as the case of model distillation, an emerging and meaningful technique that should be captured within an SBOM for AI. Indeed, this could easily lead to longer bills and redundant information, hence automation and harmonization

#

of the format is essential for creating a meaningful and effective SBOM for AI. Moreover, **it is essential to develop a framework to effectively track AI vulnerabilities and weaknesses**, given the still largely experimental results in the field of AI model red teaming.

### *Future G7 work*

This paper presented a shared G7 vision on SBOM for AI to increase transparency and cybersecurity along the full supply chain of AI systems and models. A trustworthy SBOM for AI:

- allows all the stakeholders involved in the AI supply chain to benefit from the improved transparency and knowledge of the system components;
- reduces risks, improves insight and traceability of the core components of an AI system, including security guardrails, vulnerability management and compliance attestations;
- can foster interoperability with or be integrated in already established safety, transparency and cybersecurity frameworks for traditional software, such as SBOM or security advisories and bulletins.

To fully harness the benefits and address the challenges of SBOM for AI, the next steps for the *G7 Cybersecurity Working Group – Smarter Together: Artificial Intelligence* will be to focus on providing a shared technical vision tackling these challenges, starting with a **status quo analysis of existing frameworks** to be carried out in the second half of 2025. This will be followed by further work on **technical recommendations and guidelines,** paving the way for the definition of a **common G7 framework** fostering adoption of SBOM for AI by public and private sector operators.

## References

#

Allen D. Householder, Vijay S. Sarvepalli, Jeff Havrilla, Matt Churilla, Lena Pons, Shing-hon Lau, Nathan M. VanHoudnos, Andrew Kompanek, and Lauren McIlvenny: Lessons Learned in Coordinated Disclosure for Artificial Intelligence and Machine Learning Systems. 2024.

Federal Office for Information Security (BSI): Transparency of AI Systems, White Paper. 2024.

Federal Office for Information Security (BSI): Technical Guideline TR-03183: Cyber Resilience Requirements for Manufacturers and Products - Part 2: Software Bill of Materials (SBOM). 2024.

French National Cybersecurity Authority (ANSSI): Building trust in AI through a cyber risk-based approach. Joint high-level risk analysis on AI. Version 1.0, 2025.

Ministry of Economy, Trade and Industries (METI) of Japan: Revised Guide Formulated on Specific Methods for Managing Software Vulnerability Utilizing “Software Bill of Materials (SBOM),” a List of Software Components, as a Preparatory Guide for Cyberattacks.

National Cyber Security Center (NCSC): Guidelines for secure AI system development. 2023.

The United States Department of Commerce: The Minimum Elements For a Software Bill of Materials (SBOM), 2021.#