David Soden

MARKET ASSESSMENT | AUGUST 2026

Five Acronyms, Two Schemas

What an AI bill of materials can and cannot attest, and which BOM types are real enough to put in a contract today

David Soden

Independent analyst and practitioner, enterprise automation and applied AI

davidsoden.com/market-assessment

REPORT DS-2026-1812 · 2026 · VERSION 1.0 · PUBLIC

About this report

Every substantive claim in this report carries an evidence classification and, where applicable, a numbered source. Facts are separated from vendor claims, third-party estimates, the author's assessments, and untested hypotheses. Forward-looking sections use scenarios with observable tripwires rather than point forecasts. Stated assumptions are listed before the analysis begins.

Market Assessment reports are published and commissioned at **davidsoden.com/market-assessment**. This document is analysis, not investment, legal, or procurement advice.

Scope and evidentiary standard

This assessment covers the decision facing any organisation that ships software containing AI components: whether to build a bill-of-materials programme covering models and datasets, what standard to build it against, what the resulting file can and cannot be used for, and which of the five bill-of-materials types now being marketed together are real enough to write into a contract. The five are SBOM for software, CBOM for cryptography, AIBOM for artificial intelligence, HBOM for hardware, and QBOM for quantum. They are usually presented as one framework, often under the name XBOM or xBOM.

The analysis is industry agnostic. A bank, a medical device maker, a telecommunications operator and a public sector body face the same schema questions and, with the exception of sectoral conformity assessment, the same regulatory ones. Where a vertical does change the answer, the text says so.

Two things about the evidence in this category need stating before anything else. The first is that search results on these terms are dominated by content marketing from firms that sell bill-of-materials tooling, published as neutral explainers, glossaries and academies. Several of the most confident pages on QBOM and AIBOM are sales assets with a definition wrapped around them. None of them is treated as a source here. Where a claim exists only in vendor material it is labelled VENDOR CLAIM and the vendor's product is named alongside it.

The second is that adoption figures in this category are close to worthless. There is no non-commercial count of how many organisations produce an AI bill of materials in production. The European Union Agency for Cybersecurity ran the most substantial independent adoption survey available and it does not mention AI at all. That absence is reported here as a finding rather than filled with a vendor estimate.

Claims are sourced to specification text, regulatory text, government publications and peer-reviewed research wherever such a source exists. Four primary specifications and guidance

documents totalling roughly seven hundred and fifty pages were read rather than summarised, and several of the numbers below are counts taken directly from those files. Every claim carries one of the six labels below.

Label	Definition	How to treat it
FACT	Verifiable in the cited primary source: a filing, press release, official documentation, or standards body announcement.	Reliable as stated. Check the source if the exact wording matters.
VENDOR CLAIM	Reported by a vendor about its own business or product. Self-measured and not audited by any third party.	The vendor said it. Directionally useful, not a measurement. Definitions of the metric are usually undisclosed.
THIRD-PARTY ESTIMATE	A research firm forecast or survey result. Methodology is proprietary and generally not reproducible.	Use for direction and order of magnitude only. Do not build a model on it.
ASSESSMENT	The author's reasoned judgment from the cited evidence. Falsifiable, and the reasoning is shown so it can be argued with.	This is analysis, not reporting. Disagree with it where your evidence differs.
HYPOTHESIS	A proposition offered for testing. Not currently supported by sufficient evidence to assert.	Treat as a research question, not a conclusion. Each one names what would confirm or refute it.
SCENARIO	A structured future state with a stated subjective probability. The probability is the author's judgment, not a calculated figure.	Use the leading indicators attached to each, not the probability. The indicators are observable; the probability is not.

One disclosure belongs in the scope rather than the endnotes. The framing this report examines is not neutral. "Delivery Shield" is the application security posture management platform sold by OpsMx, which markets continuous xBOM tracking spanning SBOM, CBOM and AIBOM through a single dashboard, plus a deployment bill of materials of its own devising. OpsMx sells the aggregation layer, and a five-part framework is a larger aggregation layer than a one-part framework. That commercial interest does not make the framing wrong. It does mean the count of BOM types is a thing to verify rather than accept.

Assumptions this analysis rests on

Four premises are assumed rather than demonstrated. Each is stated with the conclusions that fail if it turns out to be wrong.

First, that CycloneDX and SPDX remain the two schemas in play and no third contender emerges. If a major cloud provider or a regulator introduced a competing machine-readable format with

real distribution, the format recommendation in this report would be premature and the conversion argument would need reworking.

Second, that the European Cyber Resilience Act application date of 11 December 2027 holds. The Digital Omnibus deferred the AI Act's high-risk obligations by sixteen months in 2026, which establishes that these dates move. If the CRA slips, the only hard commercial deadline in this analysis slips with it and the case for starting now weakens considerably.

Third, that the regulatory posture of requiring documentation in general terms, rather than naming a file format, persists through the next revision cycle. Nearly every buy-side recommendation here follows from that posture. If a harmonised standard or an implementing act names a format, the calculus changes and the report should be re-read as historical.

Fourth, that model provenance metadata upstream of the enterprise does not materially improve. The completeness ceiling described later is a property of what model publishers document, not of the tools that read it. If publishers began populating provenance fields to the standard the specifications allow, the AI bill of materials would answer questions it currently cannot, and the section on its limits would be wrong.

The call

ASSESSMENT Build the AI bill of materials, build it in CycloneDX, and stop calling it an AIBOM. There are two schemas behind the five acronyms, not five artifacts, and QBOM has no specification anywhere outside one vendor category and one Indian advisory. The file is worth producing because procurement will ask for it and the Cyber Resilience Act already compels its software half. It is not worth believing: a bill of materials attests inventory, never authorship or integrity, and no quantity of additional acronyms changes that. I would move if a harmonised standard or an implementing act named a specific machine-readable format and attached a conformance test to it.

My judgment on the evidence assembled here, nothing more. There is data I can't see and data I may have missed. New evidence can move me, including to the opposite position, and I reserve the right to move.

Executive summary

A reader who has never encountered the terms AIBOM, QBOM or HBOM should know that they describe files, not systems. A bill of materials is a structured inventory of what a product is made of, expressed in a format a machine can read, so that when a defect is found in a component the organisations using that component can be identified quickly. The software version of this idea won its argument in December 2021, when a vulnerability in a logging library called Log4j sent every large enterprise hunting through its own estate for something it could not enumerate. Everything since then has been an attempt to extend that idea into new domains. The question is how far the extension actually goes.

FINDING 1 Five acronyms, two schemas, and one of the five has no specification at all

FACT CycloneDX 1.7, ratified as ECMA-424 second edition in December 2025, is a single specification with a single JSON serialisation in which hardware devices, machine learning models, datasets and cryptographic assets are four values of one component type enumeration. SPDX 3.0 reaches the same place by a different route, through profiles layered on a common core. The 566-page ECMA-424 text contains no occurrence of the string QBOM, and neither does the first cross-domain academic survey of the field, whose five searched categories are SBOM, CBOM, AIBOM, HBOM and SaaSBOM. [1] [2] [17]

FINDING 2 No binding instrument anywhere uses the words AIBOM, QBOM or HBOM

FACT The Cyber Resilience Act requires a software bill of materials in a commonly used and machine-readable format covering at least top-level dependencies, names no format, and mentions no other bill of materials type. The AI Act's Article 11 and Annex IV require technical documentation in narrative form, including datasheets on training data provenance, and use none of the acronyms. The strongest binding AI transparency obligation is Article 53(1)(d), which mandates a public prose summary on a Commission template, not a machine-readable inventory. Everything else in the five-part framework is voluntary guidance or vendor framing. [7] [8] [9]

FINDING 3 The most recent government guidance on AI bills of materials refuses to call them AIBOMs

FACT The G7 Cybersecurity Working Group's Software Bill of Materials for AI: Minimum Elements, published jointly by seven national cybersecurity agencies and the European Commission in 2026, defines seven clusters of elements across 26 pages and uses the term AIBOM zero times. So does CISA's 2026 Minimum Elements for a Software Bill of Materials, co-signed by eighteen agencies, which states in terms that it does not introduce additional elements for SBOMs for AI systems. The naming convention chosen by the agencies is SBOM for AI. [4] [5]

FINDING 4 QBOM is two incompatible definitions, neither of which is a standard

FACT India's CERT-In, the only national authority that has formalised the term, defines a QBOM as an inventory of a quantum computing device: its processor, simulators, sensors, SDKs and energy consumption. Vendors selling post-quantum readiness tooling define it as an inventory of cryptographic algorithms scored for quantum vulnerability, which is a cryptography bill of materials under a different name. CERT-In itself gives QBOM and CBOM a single shared section and a single minimum-elements table. An organisation that does not own a quantum computer has nothing to put in the CERT-In version. [11] [28]

FINDING 5 The SBOM value case does not transfer cleanly to models, and the gap is upstream

FACT The Log4j question is: which of my systems contain this component. It works because libraries are versioned, patchable and declared in manifests. Models are not patched, and the provenance an AI bill of materials would need is frequently absent at source. More than sixty percent of models and seventy percent of datasets on Hugging Face lack a complete model card. In an industrial case study, external APIs and scraping populated forty percent of AI bill-

of-materials fields for third-party models. Even ImageNet and CIFAR-10 do not fully disclose their data sources. [3]

FINDING 6 A bill of materials attests inventory and cannot attest authorship or integrity

ASSESSMENT This is the structural argument and the reason to be careful about what the artifact is asked to carry. ECMA-424 says so itself: the standard defines no enforcement mechanism for verifying the accuracy or completeness of a BOM. CISA's 2026 guidance places accuracy, coverage and completeness explicitly out of scope. The G7 authors write that an SBOM for AI by itself is not sufficient for increasing supply chain cybersecurity. Three independent bodies have documented the same limit, and the marketing around the five-part framework does not mention it. [1] [4] [5]

FINDING 7 Listing a model does not tell you the weights were not poisoned, and current scanners are close to blind

FACT Academic work published in 2025 identified 22 distinct pickle-based model loading paths across five major machine learning frameworks, of which 19 are missed entirely by existing scanners, and demonstrated 133 exploitable gadgets achieving close to a 100 percent bypass rate against available detection tools and 89 percent against the best of them. An entry in an inventory naming a model and a hash establishes which artifact you loaded. It does not establish that the artifact is what its publisher intended to ship. [18]

FINDING 8 The inventory nobody is producing is the agent tool surface

ASSESSMENT Model Context Protocol server entries and project-scoped agent configuration files execute on folder-open and sit outside the dependency graph that software composition analysis inspects. Scanners for these files now exist, from Snyk, Cisco and others, but they emit findings rather than inventory: none of them produces a bill of materials. Snyk's aibom command is the only tool identified here that places MCP clients, servers and their exposed tools into a CycloneDX document, and it does so by reading source code in four languages rather than by reading configuration files. [20] [22] [27]

FINDING 9 CycloneDX is the correct format choice, and the choice is close to reversible

ASSESSMENT CycloneDX carries the largest share of production use in the only independent adoption survey available, at 44 percent against SPDX's 29 percent. Both the OWASP AI bill-of-materials generator and Snyk's aibom command emit CycloneDX and neither emits SPDX. CycloneDX expresses AI, hardware and cryptography as component types in one document rather than requiring profile composition. Protobom and bomctl provide a conversion path that is close to lossless but not identical, so the decision is cheap to revisit and not free. [1] [16] [20] [23] [26]

FINDING 10 The recurring cost is metadata ownership, not generation, and it is identical whether you build or buy

ASSESSMENT Generating a file in continuous integration is a solved problem and costs almost nothing. Keeping it true across every fine-tune, dataset refresh and prompt version, with a named owner accountable per model, is the cost that decides whether the programme survives its second year. That cost scales with model count and is unaffected by the tooling

> decision, because no tool can discover a lineage its publisher never documented. Below roughly ten production models the correct answer is a maintained spreadsheet and no programme at all. [3] [16] [21]

How many standards are actually behind the five acronyms

Start with the specification text rather than with anybody's summary of it, because this is the question the summaries get wrong most often.

FACT CycloneDX version 1.7 was released on 21 October 2025 and adopted by the Ecma International General Assembly in December 2025 as ECMA-424, second edition. Its scope clause defines one specification covering "software and hardware components, services, dependencies, vulnerabilities, cryptographic artefacts, machine learning models, and other elements relevant to supply chain transparency and cybersecurity assurance", serialised in a single machine-readable JSON format. It is one file format, not five. [1]

FACT Inside that specification, the property at /components/[]/type is an enumeration. Among its values are device, described as "a hardware device such as a processor or chip-set"; machine-learning-model, described as "a model based on training data that can make predictions or decisions without being explicitly programmed to do so"; data, "a collection of discrete values that convey information"; and cryptographic-asset, "a cryptographic asset including algorithms, protocols, certificates, keys, tokens, and secrets". HBOM, AIBOM and CBOM are, at the level of the schema, three values of one field. [1]

FACT The specification's own framing confirms this. Under the heading "xBOM Capabilities" it lists eleven capabilities of the single standard, not five separate standards: SBOM, SaaSBOM, HBOM, ML-BOM, CBOM, OBOM for operations, MBOM for manufacturing, a bill of vulnerabilities, a vulnerability disclosure report, VEX, and CycloneDX Attestations. It states that the inventory of models and datasets "can be used independently or combined with the inventory of software and hardware components or services defined in HBOMs, SBOMs, and SaaSBOMs". These are views on one document, or separate documents in one schema, at the author's discretion. [1]

ASSESSMENT That list of eleven is worth holding next to the marketed framework of five. The five-part XBOM story is not an expansion of the standard. It is a selection from it, minus one item that is not in it, presented as a framework. A reader who accepts the framing will assume there are five specifications to comply with. There is one, and it already does more than the framing claims.

FACT SPDX arrives at the same structure by a different design. SPDX 3.0, released April 2024 with version 3.0.1 following in December 2024, replaced a monolithic document with a profile architecture: a Core profile defining foundational classes, and profiles for Software, Build, Security, Licensing, AI and Dataset layered on top. The AI profile defines three classes and eighteen properties, and its AIPackage class is a specialisation of the Software profile's package. Dataset provenance lives in the Dataset profile. These are not separate artifacts either, they are namespaces in one model. [2]

FACT The SPDX AI and Dataset profiles were the product of a multi-stakeholder effort involving more than ninety contributors, documented in a peer-reviewed experience report presented at ICSE-SEIP in April 2026. The extension added 36 new fields treating datasets, models and their provenance as first-class supply chain elements. SPDX 2.2.1 is ISO/IEC 5962:2021; the 3.0 revision is at Draft International Standard stage as ISO/IEC 5962, which means the ISO number in a contract today still points at the older document. [2] [3]

ASSESSMENT So the count is two schemas and four domains, not five artifacts. SBOM, CBOM, AIBOM and HBOM are domain vocabularies expressed on one of two competing machine-readable models. The acronyms are useful as shorthand for the question being asked of the file. They are misleading as descriptions of what has to be produced, because they imply four production pipelines where one pipeline with four component types will do.

QBOM, separately

FACT No standards body has published a QBOM format. The 566 pages of ECMA-424 second edition contain seven occurrences of the word quantum, every one of them inside the field nistQuantumSecurityLevel, which records a NIST post-quantum security strength category for a cryptographic algorithm. It is a property of a cryptographic asset, not a document type. SPDX has no quantum profile. The first cross-domain academic survey of bills of materials, published by researchers at the University of New South Wales in January 2026, searched five BOM categories across the leading software engineering and security venues. Those five are SBOM, CBOM, AIBOM, HBOM and SaaSBOM. The word QBOM does not appear in the paper. [1] [2] [17]

FACT The only formalisation is India's CERT-In, whose Technical Guidelines on SBOM, QBOM and CBOM, AIBOM and HBOM version 2.0 were dated 9 July 2025. Its QBOM is an inventory of a quantum computing device. The minimum elements are model name, defined as "a unique identifier for the quantum device or system", version, vendor and origin, licence, cryptographic asset, communication protocol, hardware including "the processor, simulators, networking components, and sensors required for quantum operations", software dependencies including SDKs, environmental impact, vulnerabilities and attestations. CERT-In gives QBOM and CBOM one shared section and one shared table. [11]

VENDOR CLAIM The definition circulating in commercial material describes something else entirely. A representative vendor explainer defines a QBOM as "a structured inventory of every cryptographic algorithm in use across a software system, with each algorithm scored for its vulnerability to quantum computing attacks". That is a cryptography bill of materials with a risk score attached. The page belongs to a firm selling post-quantum migration services, and the same definition recurs across several vendors selling the same category. [28]

ASSESSMENT Two definitions that describe different things is not a maturing standard, it is an unowned term. One inventories quantum hardware you probably do not have. The other renames an inventory that already has a specification. Neither is a schema, neither has a conformance test, and neither can be sensibly written into a contract, because the two parties would be agreeing to different deliverables. QBOM is the fifth item in the framework and it is the one to strike.

FIGURE 1 Occurrences of each acronym in the primary documents that supposedly mandate it

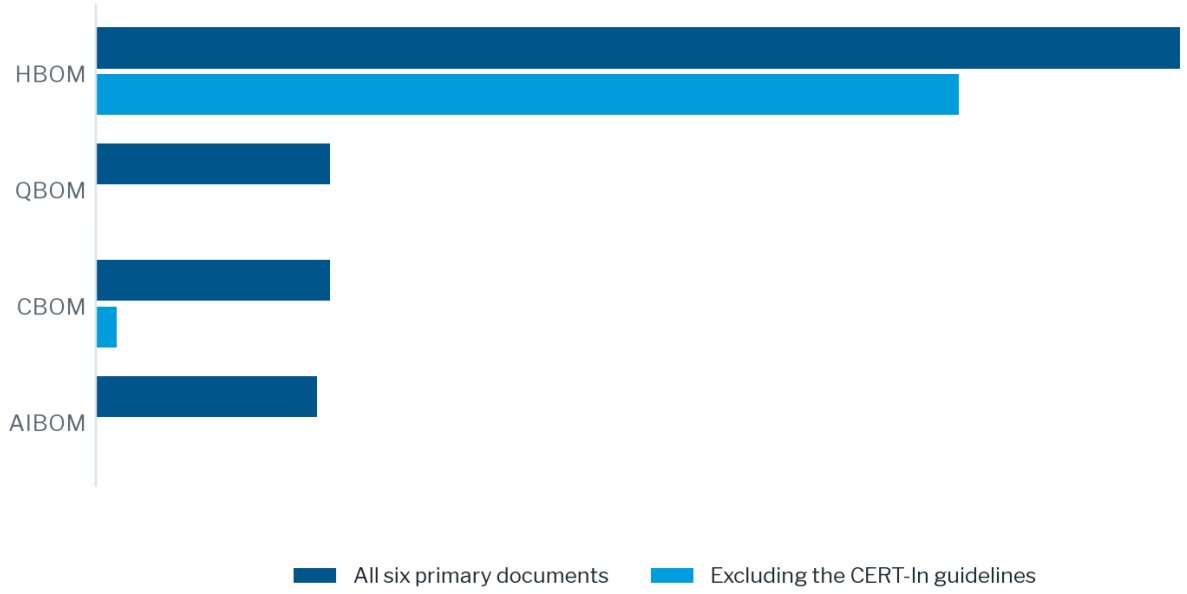

FACT Word-boundary counts taken by the author from the full text of six primary documents totalling 747 pages: ECMA-424 second edition, CISA's 2026 Minimum Elements, the G7 Software Bill of Materials for AI, CISA's 2023 HBOM Framework, CERT-In's guidelines version 2.0, and the ENISA adoption survey. SBOM is excluded from the chart because it would flatten everything else: it appears 990 times across the six, or 660 times excluding CERT-In. The HBOM total is inflated by CISA's dedicated HBOM framework, which accounts for 117 of the 125. Remove the single Indian advisory and AIBOM and QBOM disappear entirely. [1] [4] [5] [11] [12] [16]

What the regulations actually require

The marketing for the five-part framework says the transition exists to meet evolving regulatory mandates like NIST and EU guidelines. That sentence is doing a great deal of work, so it is worth separating the instruments that bind from the guidance that does not, and both from the vendor framing that maps a product onto whichever of the two sounds more urgent.

The one hard software deadline

FACT Regulation (EU) 2024/2847, the Cyber Resilience Act, entered into force on 10 December 2024. Annex I Part II point 1 requires manufacturers to "identify and document vulnerabilities and components contained in products with digital elements, including by drawing up a software bill of materials in a commonly used and machine-readable format covering at the very least the top-level dependencies of the products". Reporting obligations under Article 14 apply from 11 September 2026 and the main obligations from 11 December 2027. The SBOM does not have to be published, only retained in technical documentation and produced to market surveillance authorities on request. [7]

FACT Annex I names no format. It says commonly used and machine-readable, and no harmonised standard or delegated act specifying a format has been published. It also mentions no bill of materials other than the software one: no AI, no cryptography, no hardware. Top-level dependencies is a floor, not a ceiling, and it is a much lower floor than most tooling produces by default. [7]

What the AI Act actually says

FACT Article 11 of the AI Act requires technical documentation for high-risk systems to be drawn up before market placement and kept up to date, containing at least the elements in Annex IV. Annex IV is nine numbered headings of narrative description. Point 2(a) requires documenting "recourse to pre-trained systems or tools provided by third parties and how those were used, integrated or modified by the provider". Point 2(d) requires, where relevant, "datasheets describing the training methodologies and techniques and the training data sets used, including a general description of these data sets, information about their provenance, scope and main characteristics". Neither Article 11 nor Annex IV uses the words bill of materials, and neither names a machine-readable format. [8]

FACT Those obligations have moved. The Digital Omnibus on AI reached provisional political agreement on 7 May 2026, received final Council approval on 29 June 2026, and entered into force on 27 July 2026. It postpones the high-risk obligations for Annex III systems from 2 August 2026 to 2 December 2027, and for high-risk AI embedded in regulated products to 2 August 2028. The Article 50 transparency rules and the Article 4 AI literacy duty were left where they were. [10]

FACT The strongest binding transparency obligation in the AI Act is not in Article 11. Article 53(1)(d) requires providers of general-purpose AI models to publish a sufficiently detailed summary of training content according to a template provided by the AI Office. The Commission published that template on 24 July 2025, it applied to models placed on the market from 2 August 2025, and summaries must be updated at least every six months or on material change including fine-tuning. This is the one place a European instrument mandates a specific document form for AI provenance, and the form it mandates is a public prose summary, not a machine-readable bill of materials. [9]

FACT Compliance with that obligation is thin. An academic assessment by researchers at Trinity College Dublin's ADAPT Centre and the Mozilla Foundation, presented at ACM FAccT in 2026, conducted an exhaustive search as of 12 January 2026 and located five published summaries in total, five months after the obligation began to apply. [19]

ASSESSMENT That last number deserves more attention than it gets. The nearest thing in force to a regulatory AI provenance mandate is a mandatory template, applying to the largest model providers in the world, and an exhaustive search found five completed instances. Any procurement conversation that assumes suppliers can readily produce structured provenance for the models they ship should be tested against that figure first. [19]

The guidance that is not a mandate

FACT CISA published draft 2025 Minimum Elements for a Software Bill of Materials on 22 August 2025 under docket CISA-2025-0007, with comments closing 3 October 2025. The final 2026 Minimum Elements followed on 29 July 2026, co-signed by CISA, NSA, FBI and fifteen further national agencies including India's CERT-In, Germany's BSI and France's ANSSI. It adds ten data fields, among them an author signature, the generating tool, the data format, the lifecycle

phase, a component hash and its algorithm, and the component licence. It contains no occurrence of AIBOM or HBOM. [4] [6]

FACT On artificial intelligence, both the 2025 draft and the 2026 final say the same thing in nearly the same words: "this document does not introduce additional elements for SBOMs for AI systems". The 2026 version adds that there may be useful supply chain data of the kind AI engineers discuss under the titles model cards or data cards, and points the reader at a separate publication. [4] [6]

FACT That separate publication is Software Bill of Materials for AI: Minimum Elements, released in 2026 by Germany's BSI, Italy's ACN, France's ANSSI, Canada's CSE, the US CISA, the UK's NCSC and Japan's NCO, with the European Commission, under the G7 Cybersecurity Working Group. It defines seven clusters: Metadata, System Level Properties, Models, Dataset Properties, Infrastructure, Security Properties, and Key Performance Indicators. Its executive summary states that the elements "are not mandatory; do not create requirements, standards, or legislation". Across 26 pages it does not use the term AIBOM once. [5]

FACT CISA's Hardware Bill of Materials Framework for Supply Chain Risk Management, published September 2023 by the ICT Supply Chain Risk Management Task Force, is explicitly voluntary. It is not a format. It is a use case taxonomy and a data field taxonomy, and Appendix C maps its fields one-to-one onto CycloneDX and SPDX where a mapping exists, recording "None, do not map" where it does not. The document recommends that future work merge the HBOM framework into the emerging SBOM frameworks. [12]

FACT CERT-In's guidelines are also voluntary. They encourage departments and organisations to make SBOM creation a mandatory standard practice in their own procurement, which is a recommendation to impose a requirement rather than a requirement. On format, they specify that public sector organisations should use SPDX or CycloneDX. Even the single document that names all five acronyms points at the same two schemas. [11]

The cryptography side, which is a real deadline attached to a fictional artifact

FACT NIST published FIPS 203, 204 and 205 on 13 August 2024, standardising ML-KEM for key encapsulation, ML-DSA for digital signatures and SLH-DSA as a hash-based signature alternative. NIST IR 8547 sets the transition timeline: quantum-vulnerable public key algorithms including RSA, ECDSA, ECDH and finite-field Diffie-Hellman are deprecated after 2030 and disallowed after 2035. The NSA's CNSA 2.0 sets category-specific dates, with most transitions falling between 2030 and 2033, and excludes SLH-DSA from national security systems. [13] [14] [15]

ASSESSMENT None of those instruments requires a CBOM. They require that you know where your cryptography is, which is a genuine and difficult discovery problem, and they say nothing about the file format you record the answer in. CycloneDX happens to have a well-developed cryptographic asset model, which makes it a sensible place to put the answer. That is a tooling convenience, not a compliance obligation, and a vendor presenting the 2030 date as a CBOM deadline is compressing two different things. [1] [14]

Instrument	Status	What it actually requires	Date that matters
EU Cyber Resilience Act, Annex I Part II(1)	Binding	A software bill of materials in a commonly used, machine-readable format, top-level dependencies at minimum. No format named. No other BOM type mentioned.	11 Dec 2027 (Art. 14 reporting from 11 Sep 2026)
EU AI Act, Art. 11 and Annex IV	Binding	Narrative technical documentation including third-party pre-trained systems and training data provenance. No BOM, no format.	2 Dec 2027 for Annex III; 2 Aug 2028 embedded
EU AI Act, Art. 53(1)(d)	Binding	Public prose summary of training content on the AI Office template, refreshed at least every six months.	In force since 2 Aug 2025
CISA 2026 Minimum Elements	Voluntary guidance	Baseline SBOM fields for all software including AI. States it adds no AI-specific elements.	Published 29 Jul 2026
G7 SBOM for AI, Minimum Elements	Voluntary guidance	Seven clusters of AI supply chain elements. States it creates no requirements or standards.	Published 2026
CISA HBOM Framework	Voluntary guidance	Use case and data field taxonomy mapped onto CycloneDX and SPDX. Not a format.	Published Sep 2023
CERT-In BOM Guidelines v2.0	Voluntary guidance	Minimum elements for all five acronyms. Recommends SPDX or CycloneDX as the format.	Dated 9 Jul 2025
NIST FIPS 203/204/205, IR 8547, CNSA 2.0	Binding for US federal and NSS	Migration off quantum-vulnerable algorithms. Requires cryptographic discovery. Names no BOM.	Deprecated after 2030, disallowed after 2035

ASSESSMENT Read down the middle column and the pattern is consistent. Regulators require documentation and inventory in general terms. Vendors map those general terms onto a named artifact they sell, and the mapping is usually defensible and occasionally load-bearing in ways the regulation does not support. Exactly one binding instrument in the table requires a bill of materials at all, and it requires the software one.

What an AI bill of materials can be asked to answer

The software bill of materials earned its place by answering one question well. In December 2021 a remote code execution vulnerability in Log4j sent every large organisation searching for a component it had no inventory of. The value case is: when a CVE lands, which of my systems contain it, and how fast can I enumerate them. Test whether the AI version of that question survives contact with how models actually work.

ASSESSMENT It survives partially, and the part that survives is the part that was already an SBOM. An AI system's Python dependencies, its inference server, its CUDA libraries and its vector database client are ordinary software components with ordinary CVEs, and an ordinary SBOM finds them. That is precisely CISA's reasoning for adding nothing: since AI is software, the minimum elements illuminate many of the components. The AI-specific part is where the analogy breaks. [4]

ASSESSMENT Three properties break it. Models are not patched, so knowing you run a given model version does not point at a remediation the way a library version does; the remediation is a retrain or a swap, on someone else's schedule. A hosted inference endpoint returns a name and a version string and nothing about training data lineage, which means the most common deployment shape produces the least informative inventory entry. And the vulnerability feed that makes an SBOM operationally useful has no equivalent: there is no CVE database of model defects with the coverage and discipline of the National Vulnerability Database.

How complete are production AI bills of materials

ASSESSMENT Nobody has published a measurement of completeness for AI bills of materials produced in production environments, and this is a genuine research gap rather than an oversight in the search. What exists is a set of adjacent measurements of the upstream metadata such a file would have to draw on, and those are consistent enough to support an inference.

FACT More than sixty percent of models and seventy percent of datasets on Hugging Face lack a complete model card. In the industrial case study accompanying the SPDX AI profile work, external APIs and web scraping successfully populated forty percent of fields for third-party models, and sixty percent of existing model card fields proved extractable from the specification's own fields. The authors record that even foundational datasets including ImageNet and CIFAR-10 do not fully disclose their data sources, and that a complete AI bill of materials for a self-driving car could span several gigabytes, making generation and storage impractical. [3]

ASSESSMENT The design response to that reality is instructive. The working group deliberately optimised for adoption over completeness, defining a small set of readily recordable required fields and deferring ambitious ones, because early drafts that attempted to capture every conceivable detail met consistent practitioner pushback. Rather than enumerate every form of bias, they introduced a single knownBias field capturing only documented biases. A specification designed around what publishers will actually fill in is a specification that has already conceded the provenance question. [3]

FACT The OWASP AI bill-of-materials generator ships a completeness score for the same reason. Its configurable field registry scores 29 fields across five categories: four required fields, five metadata, five component basic, fourteen component model card and one external reference, weighted at twenty, twenty, twenty, thirty and ten percent of the total respectively, with critical fields carrying a three times multiplier. Only fields extracted with medium or high confidence contribute to the score. A tool that ships a scoring system for its own output is telling you the output is routinely incomplete. [23]

FACT The same pattern holds for software, where the tooling is fifteen years more mature. In a study of 186 open-source GitHub projects, only seven percent of the SBOM files analysed contained all the mandatory NTIA fields. An evaluation of four generation tools across 1,000 Docker images found that all four provided little or no supplier information and that none produced full coverage of package identifiers and licences. An empirical evaluation of more than 381,000 PyPI packages found only 24 percent had SPDX-compliant licence declarations and 23 percent had no licence metadata at all. A study of four tools across 7,876 GitHub repositories found systematic inconsistency between tools for the same project. [17]

FIGURE 2 **The documentation that a bill of materials has to draw on is mostly not there**

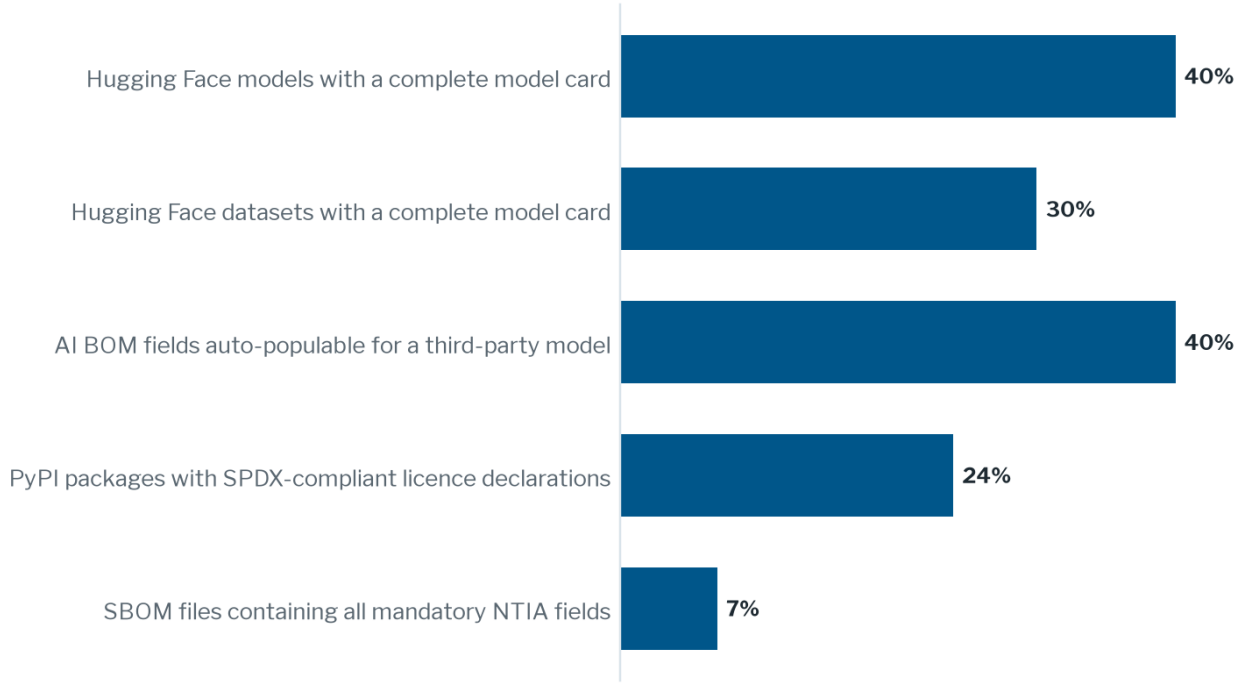

ASSESSMENT Five separate measurements from three peer-reviewed sources, combined by the author. Populations differ and the figures are not directly comparable: the first two are ceilings inferred from published statements that over 60 percent of models and over 70 percent of datasets lack a complete card; the third is from a single industrial case study; the last two are from studies of 381,000 PyPI packages and 186 GitHub projects respectively. The bars share a scale, not a methodology. The point is the direction, not the precision. [3] [17]

The fine-tune, the retrieval corpus and the prompt

ASSESSMENT A production AI system is rarely a base model. It is a base model plus a fine-tune or an adapter, plus a retrieval corpus and an embedding model, plus a system prompt, plus tool definitions, plus whatever guardrail sits in front of it. Of those, the base model is the component least likely to change and most likely to be already documented. An inventory that lists it and stops has inventoried the stable part of a system whose risk lives in the volatile part.

FACT The specifications do reach further than that criticism implies, and the gap is in practice rather than in schema. The CycloneDX AI/ML-BOM guide covers fine-tuning lineage explicitly, with worked examples declaring a base model used for fine-tuning and pointing at downstream variants including quantizations, and it covers tokenizers and prompt templates as declarable

components. The G7 minimum elements require model lineage capturing "information about the predecessor model on which fine tuning was done". Both go past the base model. [5] [24]

ASSESSMENT Retrieval is the weak point. The CycloneDX guide mentions retrieval-augmented generation twice in 91 pages and embeddings twice, against extensive treatment of training datasets. A retrieval corpus is not a training dataset: it changes daily, it frequently contains the organisation's own confidential material, and its contents determine model output more directly than the base model's weights do on any given request. There is no settled convention for expressing it as a component, and the practical result is that it is left out. [24]

ASSESSMENT Prompt versioning is the same problem with a shorter cycle. A system prompt is a configuration artifact that changes behaviour materially and is typically edited outside the software release process, sometimes through a product surface with no version control at all. CycloneDX can declare a prompt template as a component. Whether an organisation's prompt changes ever reach the pipeline that regenerates the file is an operating question, and it is the question that decides whether the artifact is true a month after it is generated. [24]

The attestation gap

This is the structural argument, and it is the reason to hold the whole category at arm's length even while producing the files.

ASSESSMENT A bill of materials attests inventory. It records what a thing is composed of, according to whoever generated the record. It does not attest that the composition is trustworthy, that the components are what their publishers intended, or that the person who authored a component is the person who was supposed to author it. Those are different claims requiring different controls, and the marketing around bills of materials consistently blurs them.

FACT The standards bodies say this plainly and the vendors do not repeat it. ECMA-424's scope clause carries an explicit note: "This Standard does not define enforcement mechanisms for verifying the accuracy or completeness of a BOM, nor does it prescribe a specific transport mechanism for exchanging BOMs." CISA's 2026 Minimum Elements places accuracy, coverage and completeness out of scope, noting they have different quality measures and assurance expectations than a signature. The G7 authors write that "an SBOM for AI by itself is not sufficient for increasing cybersecurity along the supply chain". [1] [4] [5]

FACT Academic work reaches the same conclusion from the operational side. The 2026 cross-domain survey records that for bill-of-materials data in production use "there is almost no consistent management process, no certifying authority, and no systematic quality assurance scheme", and that current frameworks rely almost entirely on static component declarations extracted from manifest files, build configurations and source code, missing runtime dependencies entirely. [17]

The precedent, from an adjacent control that already failed this way

FACT The clearest available demonstration is not from bills of materials at all. On 4 August 2026 a self-propagating worm named ChainDrop poisoned 444 npm packages and 2,212 versions in under four and a half hours. The initial compromise was a maintainer's GitHub account rather than a registry token, which meant the malicious versions shipped through npm's OIDC trusted publishing with valid SLSA provenance attestations generated by the projects' own release workflows. Provenance validation confirmed the artifact was built from the stated source by the stated workflow, which was true. Every control that validates the link between source and artifact reported success. [27]

ASSESSMENT That is a structural result, not an implementation bug. Provenance attests the path from source to artifact and cannot attest authorship, because authorship sits upstream of the point where the attestation begins. The bill-of-materials case is the same shape one step further out. An inventory attests composition and cannot attest integrity, because integrity is a property of the components rather than of the list. Adding four more domains to the list does not change what a list is capable of asserting. [27]

Applied to models

FACT An entry naming a model, a version and a hash establishes which artifact was loaded. It does not establish that the artifact is benign. Research published in 2025 mapped 22 distinct pickle-based model loading paths across five foundational machine learning frameworks and found 19 of them missed entirely by existing scanners. The same work introduced an exception-oriented bypass technique yielding nine instances, seven of which evade all scanners tested, and identified 133 exploitable gadgets achieving close to a 100 percent bypass rate overall and 89 percent against the best-performing scanner. The authors conclude that current scanners hold an incomplete understanding of the poisoning surface, leaving their detection logic fragile. [18]

FACT The G7 minimum elements do require a model hash covering "the hash of the weights or model file or other model related artifacts", which is the right field to require. They also specify what happens when it cannot be obtained: "if the SBOM author does not have access to the executable component artifact, then the SBOM author should indicate the value is unknown". For any model consumed through a hosted inference endpoint, which is the majority deployment shape, unknown is the correct and expected value. [5]

VENDOR CLAIM Set that against how the capability is sold. OpsMx, which sells the Delivery Shield application security posture management platform and is the vendor behind the five-part framing examined in this report, states that Delivery Shield "integrates AIBOM generation into your CI/CD pipeline, ensuring that every AI model deployed is transparent, accountable, and free from inherited vulnerabilities", and that aggregating SBOM, CBOM and AIBOM into a single dashboard "transforms static lists into actionable security intelligence". The claims are self-reported product marketing with no published methodology and no independent audit. [25]

ASSESSMENT Free from inherited vulnerabilities is the phrase to stop on, because it asserts the thing a bill of materials structurally cannot assert. Generating an inventory of a model establishes which model is present. Establishing that it carries no inherited defect requires detection

technology that the scanner evidence above shows does not currently exist for the model artifact itself, and requires a vulnerability feed for models that has no equivalent of the National Vulnerability Database behind it. This is not a criticism of one vendor's engineering. It is the category's standard sales register, and it is the specific claim a buyer should ask to see demonstrated before signing. [18] [25]

FIGURE 3 **What model scanners miss, from a 2025 evaluation of pickle-based poisoning**

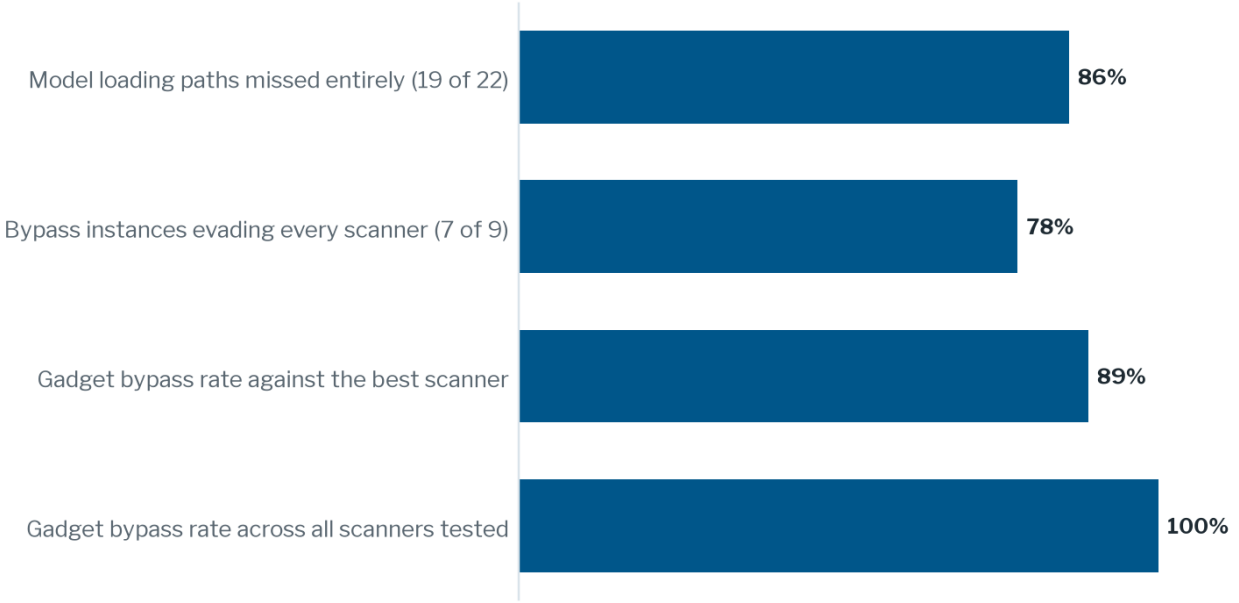

FACT Percentages derived by the author from the counts reported in the paper. The first two convert stated ratios; the second two are the paper's own rates, with the last reported as "almost a 100 percent bypass rate" and rendered here at 100. This is a single academic evaluation and the scanner set may have improved since publication. It is offered as evidence that the detection problem is open, not as a current measurement of any specific product. [18]

The inventory nobody is producing

ASSESSMENT While five BOM types are being marketed, the artifact class with the clearest and most immediate exposure has no bill of materials at all. Model Context Protocol server entries, project-scoped agent configuration files and editor task definitions execute on folder-open or on agent session start. They declare execution rather than dependencies, which places them outside the dependency graph software composition analysis inspects, which is why a clean composition report and a fully compromised checkout are compatible states. ChainDrop used exactly this surface, planting agent hooks and editor tasks across as many as fifty branches per repository. [27]

FACT Scanners for this surface now exist and they are recent. Snyk's agent-scan auto-discovers configurations for more than a dozen agent runtimes including Claude Code, Cursor, VS Code, Windsurf, Gemini CLI and Codex, scanning MCP configuration files and skill files across system, user, project and extension scopes, and scores fifteen distinct risks including prompt injection, tool poisoning and tool shadowing. Cisco publishes an MCP scanner supporting offline scanning of pre-generated MCP JSON in continuous integration. Both emit findings reports. Neither emits a bill of materials, and Snyk's own documentation describes its command-line output as experimental and subject to change. [22]

FACT Snyk's aibom command is the exception and it does what the brief on this question hoped to find. Its documentation states it identifies AI models, datasets and the AI supply chain including connections to external tools and services reached through the Model Context Protocol, categorising MCP clients, MCP servers, and the specific tools and resources an MCP server makes available. Output is CycloneDX 1.6 JSON. The stated coverage is local projects written in Python, Java, JavaScript or Go, with agents identified from popular agent libraries and tools from popular tool-calling patterns. [20]

ASSESSMENT The limitation is the interesting part. Reading source code finds the MCP connections a developer wrote in code. It does not find the ones declared in a configuration file that the agent runtime reads at startup, which is where most of them live and which is the vector ChainDrop actually used. The correct read is that one vendor has done real work here and the coverage model does not yet match the exposure. Neither CycloneDX nor SPDX has an MCP construct; Snyk maps it onto generic component and dependency structures, which works and is not standardised. [1] [2] [20]

Artifact	What it attests	What it does not, and what would
SBOM	Which components the generator observed in the build, at the time it ran	That the list is complete or accurate. ECMA-424 defines no enforcement mechanism. Binary analysis or reconciliation against the registry tarball would.
SLSA provenance	That the artifact was built from the stated source by the stated workflow	That the source was authored by the expected person. ChainDrop shipped valid provenance. Minimum release age and maintainer account controls would.
AI BOM model entry	Which model name, version and where available which weight hash was loaded	That the weights are unpoisoned. Model scanners currently miss most loading paths. Signed weights from the publisher and a reproducible fine-tune record would.
AI BOM dataset entry	What the publisher documented about the training corpus	Anything the publisher did not document, which for most models is most of it. Nothing available to the consumer closes this.
CBOM cryptographic asset	Which algorithms the discovery tool found in the code it examined	Cryptography in firmware, in third-party binaries, or negotiated at runtime. Runtime protocol observation would.
Agent configuration files	Nothing. No BOM type currently inventories them	Everything. File integrity monitoring on those paths, change review, and scanning all branches rather than the default one.

Format choice, decided

The instinct in this category is to say support both. That is the wrong answer for anyone who has to run the pipeline, because two formats means two generators, two validators, two sets of field mappings and two places for the inventory to be subtly different. Pick one.

ASSESSMENT Build against CycloneDX. Four reasons, in order of weight.

ASSESSMENT First, the tooling that exists for the AI case emits CycloneDX and does not emit SPDX. Snyk's aibom command outputs CycloneDX 1.6 JSON. The OWASP AI bill-of-materials generator outputs CycloneDX 1.6 JSON, and its documentation records that it aligns fully with the SPDX 3.0 AI profile at field level while noting it does not currently generate SPDX output and that dual-format support is future work. Choosing SPDX today means choosing the format the AI-specific tools do not write. [20] [23]

ASSESSMENT Second, the modelling is simpler for the multi-domain case, which is the case this whole framework is about. In CycloneDX a hardware device, a machine learning model, a dataset and a cryptographic asset are four values of one component type field in one document. In SPDX the equivalent requires composing the Software, AI and Dataset profiles, which is architecturally cleaner and operationally more work. If the goal is one artifact covering a system that contains models and cryptography and hardware, one component list beats three profiles. [1] [2]

THIRD-PARTY ESTIMATE Third, production use runs in its favour. In ENISA's 2026 adoption survey, 44 percent of respondents reported using CycloneDX against 29 percent using SPDX, with 17 percent on a proprietary format and 11 percent using no standard format at all. The survey drew 334 responses fielded at the end of 2025, is self-selected toward organisations already engaged with the Cyber Resilience Act, and skews European and large-enterprise. It is the best independent adoption data available and it should still be discounted for those biases. [16]

FACT Fourth, the ISO number that a procurement lawyer will reach for currently points at the older SPDX. ISO/IEC 5962:2021 is SPDX 2.2.1, which predates the AI and Dataset profiles entirely. The 3.0 revision is at Draft International Standard stage. CycloneDX 1.7 is ECMA-424 second edition, adopted December 2025, and covers the AI case in its current published form. [1] [2]

FIGURE 4 **Which format organisations report actually using**

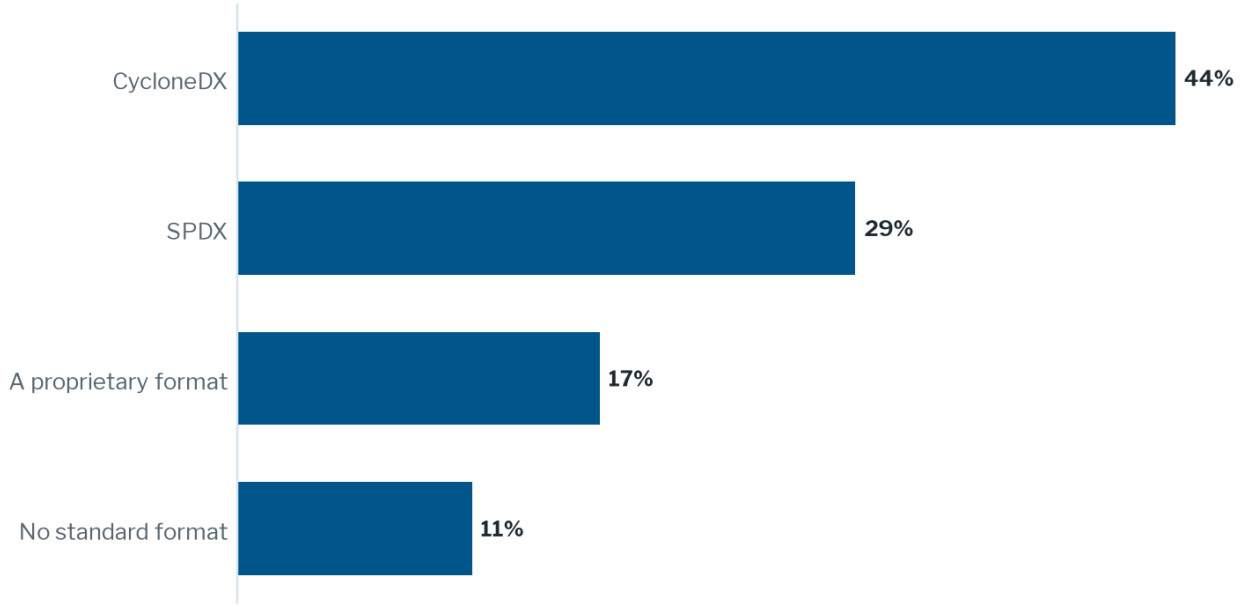

THIRD-PARTY ESTIMATE ENISA SBOM Adoption State of Play, published June 2026 from a survey fielded at the end of 2025. 334 responses, 87.7 percent private sector, more than 65 percent organisations over 250 employees, 65 percent operating mainly in the EU, more than 80 percent directly impacted by the Cyber Resilience Act. Self-selected and self-reported, with no independent verification of what respondents actually generate. Percentages sum to 101 through rounding. The 28 percent using something proprietary or nothing standardised is the more interesting number. [16]

Is the choice reversible

ASSESSMENT Mostly, and that matters more than the choice itself. Protobom, an OpenSSF project, provides a format-agnostic protocol buffers representation designed to ingest modern SPDX and CycloneDX documents and convert between them. bomctl builds a command-line tool on top of it for pipeline use. The published caveat is that conversion is close to lossless rather than lossless: tool-specific fine-grained properties can be dropped, and even a round trip with no operations does not reproduce the input byte for byte. [26]

ASSESSMENT So the decision is cheap to revisit for the core inventory and expensive to revisit for anything expressed in format-specific extensions. The practical instruction that follows is to keep the AI-specific metadata inside the standardised fields rather than in CycloneDX properties or SPDX annotations, because the standardised fields survive conversion and the extension mechanisms are exactly what conversion drops. Treat the property bag as convenience, never as the record. [26]

What the programme costs, and who ends up paying

Generating the file is not the cost. Every build-versus-buy conversation in this category starts in the wrong place because the generation step is nearly free and highly visible, while the ownership step is expensive and invisible until year two.

ASSESSMENT The recurring obligation is this: every model, fine-tune, dataset refresh, retrieval corpus change and prompt version has to reach the artifact, and somebody has to be accountable

by name for each model when the artifact is wrong. That is a metadata governance job, not an engineering job, and it is the part that decides whether the programme is still running when the person who built it changes teams.

THIRD-PARTY ESTIMATE The cost model below uses a fully loaded annual cost of 220,000 US dollars for a senior application security or platform engineer, which is the levels.fyi average total compensation figure for the role and sits at the upper end of a range where Indeed reports about 150,000 and Glassdoor about 165,000 for base salary. Those sources measure different things and none is audited. A reader in a lower-cost market should scale the whole model rather than adjust individual lines. [21]

ASSESSMENT Build has a fixed component of roughly half a full-time engineer, or about 110,000 dollars a year, covering generator integration in continuous integration, schema version upgrades, validation, storage and retrieval, and the internal support burden. Buy has a licence component plus a smaller integration component. Snyk publishes 25 dollars per contributing developer per month for its Team tier, so a 200-developer organisation is at roughly 60,000 dollars a year, with the aibom command included in the platform rather than sold separately; add about a tenth of an engineer for integration and the fixed side is around 75,000. [20] [21]

ASSESSMENT Both then carry the same variable term. Assume a named owner spends roughly four working days a year per production model on refresh and review across fine-tunes, dataset changes and prompt versions, which at a fully loaded day rate of about 850 dollars is 3,400 dollars per model per year. That term is identical on both sides, because no tool can discover a lineage the publisher never documented and no tool can decide who is accountable.

FIGURE 5 Annual cost of an AI bill-of-materials programme, by production model count

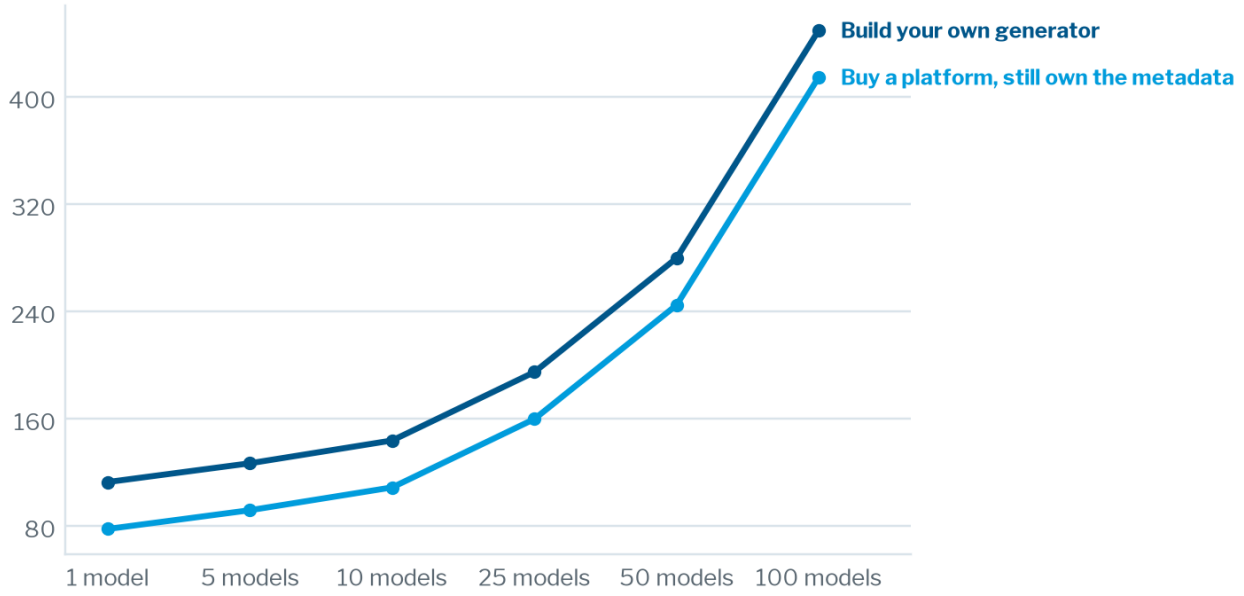

ASSESSMENT The author's model, not an observed cost. Build assumes 0.5 fully loaded engineer at 220,000 USD; buy assumes a 200-developer organisation at Snyk's published 25 USD per contributing developer per month plus 0.1 engineer for integration. Both carry an identical 3,400 USD per model per year for named-owner metadata upkeep, derived from four days at an 850 USD fully loaded day rate. The lines never cross, which is the finding: the tooling decision moves the intercept and

the ownership cost sets the slope. Sensitive to salary assumptions and to the four-day estimate, which is the softest input. [20] [21]

ASSESSMENT Buy wins on the fixed component at almost every size, and the gap is roughly 35,000 dollars a year, which is not a difficult decision for anyone already carrying a software composition analysis subscription. The reason to look at the chart is the slope. Both lines rise at the same rate, so an organisation that buys the platform and does not staff the ownership has bought the cheap half of the problem and will discover in year two that its artifacts describe a system it no longer runs.

ASSESSMENT The threshold question deserves a direct answer. Below roughly ten production models, do not build a programme. Maintain a spreadsheet with an owner per model, generate a CycloneDX file from the pipeline because it costs almost nothing to switch on, and revisit when the count doubles. The fixed cost of a real programme is roughly 75,000 to 110,000 dollars a year and it buys automation of a task that ten rows do not need. Spending it earlier is how governance programmes acquire the reputation that makes the next one harder to fund.

THIRD-PARTY ESTIMATE The barrier data supports this ordering. In the ENISA survey, the technical challenge of achieving a high degree of completeness was cited by 62 percent of respondents, data quality by 37 percent, vulnerability matching by 35 percent as quite challenging with a further 23 percent calling it extremely challenging, and lack of internal skills or dedicated staff by 28 percent. None of those is a generation problem and none of them is solved by buying a generator. [16]

FIGURE 6 **Where SBOM adoption actually sits, fifteen years into the idea**

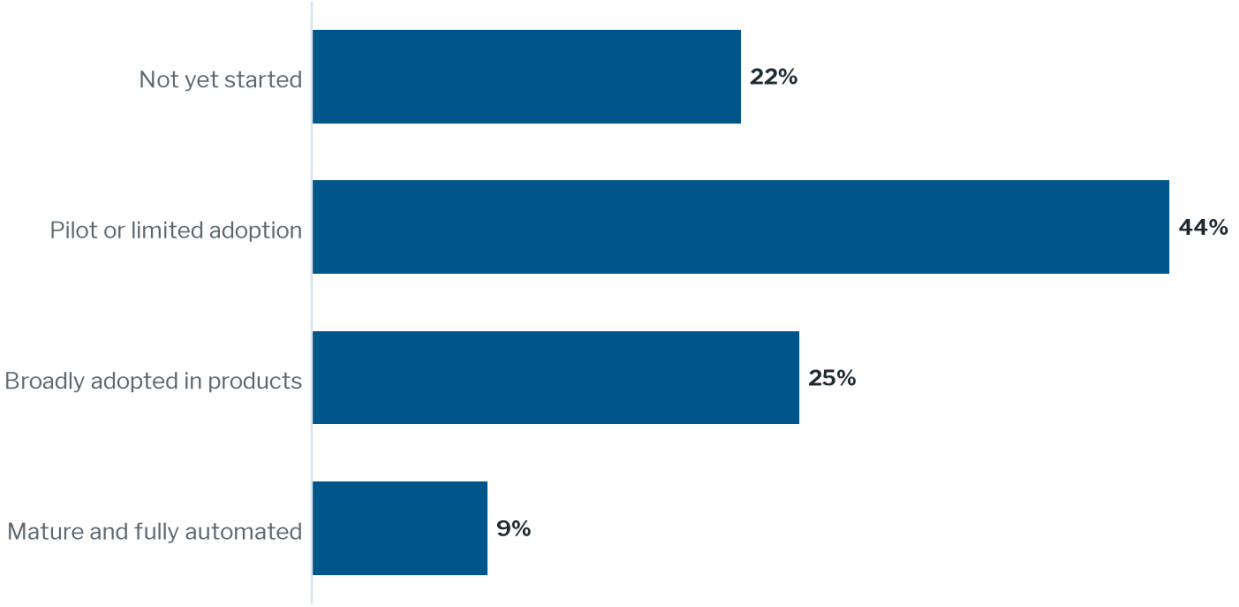

THIRD-PARTY ESTIMATE ENISA SBOM Adoption State of Play, June 2026, 334 self-selected respondents skewed toward large EU organisations already engaged with the Cyber Resilience Act. The 'not yet started' figure is derived by subtraction from the reported 78 percent who have initiated adoption. Nine percent mature after fifteen years of SBOM tooling is the relevant baseline for anyone forecasting how quickly four further BOM types will arrive. This survey does not mention AI, models or machine learning anywhere in its 33 pages. [16]

Scenarios to 2029

Probabilities here are subjective and are the least reliable content in this report. The tripwires are the useful part, because each is a public event that any reader can watch for without private information.

SCENARIO A Two schemas, more profiles 45 percent

CycloneDX and SPDX absorb each new domain as component types or profiles rather than spawning new formats. The acronyms persist as vocabulary. QBOM never acquires a specification and quietly merges into CBOM.

Consequence. The five-part framework stops being a procurement structure and becomes a way of describing sections of one file. Organisations that built one pipeline against one schema are correct and organisations that built five are carrying redundant machinery.

Earliest visible sign. A CycloneDX 2.0 or SPDX 3.2 release that adds a domain, such as agent tooling or retrieval corpora, as component types or a profile rather than as a separate specification.

SCENARIO B The harmonised standard names a format 25 percent

CEN and CENELEC publish harmonised standards under the Cyber Resilience Act that cite ECMA-424 or ISO/IEC 5962 by number, or an AI Act implementing act does the equivalent. Format choice stops being a choice for anyone selling into the European market.

Consequence. The recommendation in this report becomes either a lucky guess or an expensive migration. Conversion tooling moves from convenience to critical path, and the fidelity limits of Protobom become a compliance question rather than an engineering one.

Earliest visible sign. Publication in the Official Journal of a harmonised standard reference under the CRA that names a BOM specification by its standard number, or a delegated act on SBOM format under Article 13.

SCENARIO C Acronym proliferation continues 20 percent

More national authorities copy CERT-In's five-part structure into their own guidance. Procurement language fragments across jurisdictions and vendors package per-acronym modules against it.

Consequence. Compliance cost rises without a corresponding rise in assurance, because the underlying schemas have not changed. The organisations that suffer most are those selling into many jurisdictions with small compliance teams.

Earliest visible sign. A second national cybersecurity authority publishing guidance that names QBOM as a distinct artifact with its own minimum elements table, separate from CBOM.

SCENARIO D Attestation becomes the requirement 10 percent

> The binding obligation moves from producing an inventory to producing a verified one. Author signatures, conformance tests and third-party validation become the thing being asked for, and an unsigned or unverified BOM stops satisfying anybody.
>
> **Consequence.** The whole category re-prices. The generation tools that dominate today become commodity inputs to an assurance layer, and the organisations that treated the file as a compliance checkbox have to rebuild around evidence.
>
> **Earliest visible sign.** CISA or an EU body publishing a conformance test suite or a validator with pass criteria, rather than a minimum elements list. CISA's 2026 addition of an author signature element is the first step on this path.

ASSESSMENT Two tripwires matter more than the others and are worth stating in isolation. For QBOM to become a real format, one of three things has to appear: a working group at Ecma TC54 or in the SPDX project opening a quantum profile with published draft text; a standards body such as NIST or ETSI publishing a QBOM data field taxonomy on the model of CISA's HBOM framework; or a CERT-In version 3.0 that separates QBOM from CBOM into its own section with its own minimum elements. Until one of those happens, QBOM is a marketing term with two contradictory definitions.

ASSESSMENT For a regulator to name a BOM format by name rather than requiring documentation in general terms, watch the harmonised standards pipeline rather than the legislative one. The Cyber Resilience Act deliberately left format to standardisation, which means the naming, if it comes, arrives as a citation in the Official Journal rather than as an amendment to the regulation. That is a quieter event than a legislative change and it will be reported later, which makes it worth monitoring directly. [7]

Contract language: what to require now, what to wait on

The following is the short list this analysis produces for anyone writing bill-of-materials requirements into a supplier agreement in the next twelve months.

Acronym	What it is	State of the standard behind it	Contract today?
SBOM	Inventory of software components and their dependency relationships	Two mature schemas. ECMA-424 second edition and ISO/IEC 5962. A binding CRA obligation from 11 Dec 2027 and eighteen-agency minimum elements guidance.	Yes. Require it, with the format named.
CBOM	Inventory of cryptographic algorithms, keys, certificates and protocols	A component type in CycloneDX with a developed model. No binding obligation to produce the artifact, but a hard migration deadline that requires the underlying discovery.	Yes, as a discovery deliverable. Do not name it CBOM in the clause.

Acronym	What it is	State of the standard behind it	Contract today?
AIBOM	Inventory of models, datasets and their provenance	Real schema support in both CycloneDX and SPDX 3.0. Voluntary G7 minimum elements. No binding instrument uses the term. Upstream metadata is mostly absent.	Yes, but require named fields, not the acronym.
HBOM	Inventory of physical components in a hardware product	A component type in CycloneDX. A voluntary CISA framework from 2023 that maps onto CycloneDX and SPDX rather than defining a format.	Only if you buy hardware. Cite the CISA taxonomy.
QBOM	Contested. Either an inventory of quantum computing hardware or a renamed CBOM.	No specification from any standards body. Two incompatible definitions in circulation. One national advisory that gives it a shared table with CBOM.	No. Strike it.

ASSESSMENT Three drafting instructions follow from that table. Require fields, not acronyms: a clause that says the supplier shall provide an AIBOM is unenforceable because the parties have no shared definition, whereas a clause listing the G7 model cluster elements and naming CycloneDX 1.7 is enforceable and testable. Require the completeness score alongside the artifact, because the gap list is the governance information and a file with empty provenance fields looks identical to a file with populated ones until somebody measures it. And require regeneration on a trigger rather than on a calendar: on model change, fine-tune, dataset refresh or prompt version, because a quarterly file for a system that changes weekly is a compliance artifact rather than an inventory. [5] [23]

ASSESSMENT One thing to explicitly not require: a supplier warranty that models are free of poisoned weights or inherited vulnerabilities. No supplier can support that warranty on current detection technology, the scanner evidence above shows why, and asking for it produces either a refusal that delays the deal or an acceptance that is worthless. Ask instead for the weight hash where the supplier controls the artifact, an unknown declaration where it does not, and notification on model change. [5] [18]

Implications

For application security leads and platform engineering

ASSESSMENT You are the people who will be asked to produce these files, and the useful framing is that you are being asked to add component types to a pipeline you already have, not to build four new ones. Turn on CycloneDX generation, add the AI component types, and treat the AI bill of materials as the same artifact with more rows. Then spend the effort you saved on the thing that is not in any of the five acronyms: inventory the agent configuration files and MCP server entries in your repositories, put them under change control, and scan every branch rather than

the default one. That surface executes on folder-open, no bill of materials covers it, and the scanners that do exist emit findings rather than inventory. [20] [22]

ASSESSMENT The second instruction is about honesty in the artifact. Where a provenance field is unknown, emit unknown rather than omitting the field, which is what the G7 guidance specifies and what the OWASP completeness score rewards. An inventory that distinguishes not applicable from not obtainable from not yet done is an operational document. One that silently omits is a document that will be read as complete by somebody who was not in the room when it was generated. [5] [23]

For procurement and third-party risk

ASSESSMENT Your bargaining position is strongest at contract and your exposure is worst at the point where a clause proves untestable. Use the table above: require the software bill of materials with a named format and a named minimum element set, require AI provenance as enumerated fields rather than as an acronym, and strike QBOM. When a supplier offers an AIBOM, ask two questions before accepting it as evidence: which schema and version, and what completeness score, because the answer to the second question is the one that tells you whether the file means anything. [5] [23]

ASSESSMENT Calibrate your expectations against the Article 53 evidence. The largest model providers in the world are subject to a mandatory Commission template for training content summaries, and an exhaustive academic search a year ago found five completed instances. A mid-market supplier asked to produce structured training data provenance for a model it consumes through an API cannot do it, and a clause requiring it will produce an exception rather than a document. Write the clause to require what the supplier controls and to require disclosure of what it does not. [19]

For compliance and regulatory affairs

ASSESSMENT The mapping you are being asked to make does not exist yet, and saying so is the correct professional answer. There is no binding instrument that requires an AIBOM, a QBOM or an HBOM. The Cyber Resilience Act requires a software bill of materials and names no format. The AI Act requires narrative technical documentation and one public prose summary on a Commission template. Producing a CycloneDX file with AI component types is a reasonable way to generate raw material for Annex IV point 2(a) and 2(d), and it does not discharge either obligation, because Annex IV asks for description and rationale that no schema field carries. [7] [8] [9]

ASSESSMENT Two dates to hold. The AI Act high-risk obligations moved to 2 December 2027 for Annex III systems and 2 August 2028 for those embedded in regulated products, which bought sixteen months and demonstrates these dates are movable. The Cyber Resilience Act date of 11 December 2027 has not moved and is the one to plan against. If your organisation is treating August 2026 as the AI Act deadline, that is out of date. [7] [10]

For CISOs and engineering VPs approving the spend

ASSESSMENT The thing you are buying is inventory, and the thing you will be told you are buying is assurance. Hold the distinction, because the failure mode of this category is a board slide reporting five BOM types in green while the actual assurance question is untouched. Three standards bodies, two peer-reviewed papers and your own supplier's specification all say a bill of materials cannot attest accuracy, completeness or integrity. Anybody selling the five-part framework as risk reduction should be asked which of those three it attests, and the honest answer is none. [1] [4] [5] [17]

ASSESSMENT On the spend itself: the licence is the cheap part and buying is the right call at almost any scale, because the fixed cost of building is higher than the fixed cost of buying and neither touches the recurring cost. Fund the named owner per model or do not fund the programme. A tool with no accountable owner produces a file that is true on the day it is generated and decays from there, and the decay is invisible until an incident or an audit exposes it. If you have fewer than ten production models, decline the programme and revisit at twenty. [16] [21]

ASSESSMENT Finally, the control that is actually missing is cheap. File integrity monitoring on agent configuration paths across developer endpoints and repositories, change review on those files, and reconciliation of published artifacts against source, are a sprint of platform work rather than a product category. They close a gap that ChainDrop demonstrated at scale and that no BOM type on any vendor's roadmap addresses. [27]

What this document cannot answer

Four questions need primary research that does not currently exist, and each would materially change the analysis above.

How complete are AI bills of materials produced in production. Every completeness measurement cited here is of upstream model cards, of open-source SBOMs, or of a single industrial case study. Nobody has taken a corpus of AI bills of materials generated by enterprises against real systems and measured field population against the G7 clusters or the OWASP registry. That study would settle whether the artifact is thin because the specification is young or because the metadata does not exist, and those imply different remedies.

How many organisations produce one at all. There is no non-commercial count. ENISA's survey, the most substantial independent adoption instrument in the field, does not mention AI, models or machine learning anywhere in its 33 pages. Every adoption figure in circulation traces to a vendor survey with an undisclosed instrument, and this report declines to repeat any of them.

How many repositories carry agent configuration files that their owners did not author. ChainDrop planted across as many as fifty branches per repository, infected packages were yanked, and the configuration files persist in git history and in every clone taken during the window. A survey of public repositories would produce a real number and nobody has published one.

Whether conversion between CycloneDX and SPDX is lossless for the AI fields specifically. Protobom's published caveats concern fine-grained tool-specific properties in the general case. Nobody has round-tripped a corpus of AI bills of materials through the conversion and diffed the model and dataset fields. That measurement decides whether the format recommendation in this report is genuinely reversible or only reversible for the software half.

Half-life of this assessment

Decaying within six months: the tooling claims. Snyk's aibom coverage of MCP, its language support, and its command-line output are described in its own documentation as subject to change, and the agent-scan risk list changed between two minor versions. The OWASP generator's 29-field registry is explicitly configurable and its dual-format support is stated as planned. Treat every specific tool capability here as a snapshot.

Decaying within twelve months: the regulatory dates and the count of what is voluntary. The Digital Omnibus moved the AI Act high-risk dates by sixteen months in a single instrument, which means any date in this report can move. The harmonised standards pipeline under the Cyber Resilience Act is active and a format citation could land inside this window, which is scenario B and would revise the format section. CISA's addition of an author signature element suggests the conformance question is opening.

Decaying within twenty-four months: the QBOM finding. It is currently true that no standards body has published a QBOM format and that the term has two incompatible definitions. That is a statement about the present state of a young term and it is the single claim here most likely to be overtaken. The tripwires above are specifically designed so that a reader can check it independently rather than trusting this document.

Architectural, and unlikely to decay: that a bill of materials attests inventory and cannot attest authorship, accuracy or integrity, which is a property of what a list is rather than of any implementation, and which three standards bodies state in their own text. That the number of distinct schemas is smaller than the number of acronyms, because domain vocabularies are cheaper to invent than data models. That the recurring cost of a metadata programme is ownership rather than generation, and therefore that the tooling decision cannot fix it. And that when the framing of a technical category comes from firms selling the aggregation layer, the count of things to aggregate is a claim to verify rather than a fact to accept.

Limitations

The author was invited to appear as a panellist on a vendor-hosted webinar concerning the transition from SBOM to a five-part XBOM framework, hosted by the vendor named in the scope section. No fee was offered or paid. The vendor did not commission, fund, review or see this report before publication, has no editorial control over it, and is a subject of the analysis rather than a party to it. No other commercial relationship exists with any vendor named here, and no vendor was given advance sight.

The evidence base has four specific weaknesses worth naming. First, the acronym counts in Figure 1 are the author's own text searches across six documents, performed on extracted PDF text. Extraction can drop characters at line breaks and in tables, so treat the counts as accurate to within a few percent rather than exact. The zero results were spot-checked by hand; the large positive counts were not.

Second, the completeness figures in Figure 2 combine five measurements from three studies with different populations, methods and dates. They share a chart axis and nothing else. They support a direction and cannot support a comparison between bars.

Third, the cost model in Figure 4 is constructed rather than observed. Its softest input is the four days per model per year of ownership overhead, which is the author's estimate from adjacent metadata governance work and has no published basis. A reader whose models change monthly rather than quarterly should double it, and the conclusion about the slope survives either way.

Fourth, the only independent adoption data available is a self-selected survey of 334 respondents skewed toward large European organisations already engaged with the Cyber Resilience Act, which is close to a best case for adoption. Actual population figures are almost certainly lower than the ones reported here, and no independent AI-specific adoption data exists at all.

One further limitation of scope. This report reads specifications, regulations and published research. It does not test the tools. Claims about what Snyk's aibom command, the OWASP generator, Protobom or bomctl do are taken from their own documentation and are labelled accordingly. A hands-on evaluation would be a different and complementary piece of work.

Sources

[1] Ecma International. "ECMA-424, 2nd edition: CycloneDX Bill of materials specification." December 2025. 566 pages, read in full. *Standards body primary* ecma-international.org

[2] SPDX Project, The Linux Foundation. "SPDX Specification 3.0.1, AI Profile." December 2024. *Standards body primary* spdx.github.io

[3] Rajbahadur, G. K., Gallaba, K., Rashno, E., Suriyawongkul, A., Bennet, K., Stewart, K., Hassan, A. E. "Building an Open AIBOM Standard in the Wild: An Experience Report on Extending the SPDX SBOM (ISO/IEC 5962:2021) for AI Supply Chains." ICSE-SEIP '26, Rio de Janeiro, April 2026. *Peer-reviewed academic* arxiv.org

[4] CISA, NSA, FBI and fifteen international partner agencies. "2026 Minimum Elements for a Software Bill of Materials (SBOM)." 29 July 2026. *Government primary* cisa.gov

[5] BSI (Germany), ACN (Italy), ANSSI (France), CSE (Canada), CISA (US), NCSC (UK), NCO (Japan), with the European Commission. "Software Bill of Materials for AI: Minimum Elements." G7 Cybersecurity Working Group, 2026. 26 pages, read in full. *Government primary* bsi.bund.de

[6] CISA. "2025 Minimum Elements for a Software Bill of Materials (SBOM)," draft for public comment, 22 August 2025, docket CISA-2025-0007. *Government primary* cisa.gov

[7] Regulation (EU) 2024/2847 of the European Parliament and of the Council (Cyber Resilience Act), Annex I Part II point 1 and Article 71. *Regulatory primary* eur-lex.europa.eu

[8] Regulation (EU) 2024/1689 (Artificial Intelligence Act), Article 11 and Annex IV. *Regulatory primary* artificialintelligenceact.eu

[9] European Commission AI Office. "Template for the public summary of training content for general-purpose AI models," under AI Act Article 53(1)(d). Published 24 July 2025. *Regulatory primary* digital-strategy.ec.europa.eu

[10] Council of the European Union. "Artificial intelligence: Council gives final green light to simplify and streamline rules." Press release, 29 June 2026. Digital Omnibus on AI in force 27 July 2026. *Regulatory primary* consilium.europa.eu

[11] Indian Computer Emergency Response Team (CERT-In), Ministry of Electronics and Information Technology. "Technical Guidelines on SBOM, QBOM & CBOM, AIBOM and HBOM," Version 2.0, dated 9 July 2025. 66 pages, read in full. *Government primary* cert-in.org.in

[12] CISA ICT Supply Chain Risk Management Task Force. "A Hardware Bill of Materials (HBOM) Framework for Supply Chain Risk Management." September 2023. *Government primary* cisa.gov

[13] NIST. FIPS 203 (ML-KEM), FIPS 204 (ML-DSA) and FIPS 205 (SLH-DSA). Published 13 August 2024. *Standards body primary* csrc.nist.gov

[14] NIST. "NIST IR 8547, Transition to Post-Quantum Cryptography Standards." Deprecation after 2030, disallowance after 2035. *Standards body primary* nvlpubs.nist.gov

[15] National Security Agency. "Commercial National Security Algorithm Suite 2.0 (CNSA 2.0)." Category-specific transition dates between 2030 and 2033. *Government primary* nsa.gov

[16] ENISA. "SBOM Adoption State of Play: 2026 Survey Results and Analysis." June 2026. 334 self-selected respondents, fielded end of 2025. *Government agency survey (self-selected, self-reported)* enisa.europa.eu

[17] Zhang, S., Lyu, M., Habibi Gharakheili, H. "A Survey on Mapping Digital Systems with Bill of Materials: Development, Practices, and Challenges." University of New South Wales, January 2026. Includes the Nocera, Torres-Arias, Scarlato and Yu results cited in the text. *Peer-reviewed academic* arxiv.org

[18] "The Art of Hide and Seek: Making Pickle-Based Model Supply Chain Poisoning Stealthy Again." arXiv preprint 2508.19774, 2025. *Academic preprint (not peer-reviewed at time of writing)* arxiv.org

[19] Blankvoort, D. A. H., Pandit, H. J., Gahntz, M. "Quality Assessment of Public Summary of Training Content for GPAI models required by AI Act Article 53(1)(d)." ADAPT Centre, Trinity College Dublin and Mozilla Foundation. Submitted February 2026, presented at ACM FAccT 2026. Exhaustive search as of 12 January 2026 located five published summaries. *Peer-reviewed academic* arxiv.org

[20] Snyk. "AI-BOM," Snyk CLI command documentation. *Vendor primary (commercial interest)* docs.snyk.io

[21] Snyk. "Plans and Pricing." Team tier at 25 USD per contributing developer per month. Salary inputs from levels.fyi, Indeed and Glassdoor, August 2026. *Vendor primary and compensation aggregators (self-reported)* snyk.io

[22] Snyk. "agent-scan: Security scanner for AI agents, MCP servers and agent skills." GitHub repository. *Vendor primary (commercial interest)* github.com

[23] OWASP GenAI Security Project. "AIBOM Generator." Field registry and AI SBOM API documentation. *Open source project primary* github.com

[24] OWASP CycloneDX. "Authoritative Guide to AI/ML-BOM." 91 pages. *Standards body primary* cyclonedx.org

[25] OpsMx. "Delivery Shield" product documentation and secure software delivery pages, describing single-pane management of SBOM, CBOM and AIBOM. *Vendor primary (commercial interest, subject of analysis)* docs.opsmx.com

[26] OpenSSF. Protobom, "A universal SBOM representation in protocol buffers," and bomctl. Conversion fidelity caveats from project documentation. *Open source project primary* github.com

[27] Soden, D. "Agent Speed and the Detection Gap: Autonomous AI agents as an offensive cyber capability, and whether enterprise detection can see them." August 2026. ChainDrop reconstruction from StepSecurity, Unit 42 and The Register. *Author's prior work* davidsoden.com

[28] PostQuantum.com. "What Is a QBOM? Quantum Bill of Materials vs CBOM Explained." Representative of the vendor definition in circulation. Published by a firm selling post-quantum migration services. *Vendor-published explainer (commercial interest)* postquantum.com

Rights and permitted use

It is published for public reading. You are free to share the complete, unmodified file, to link to it, and to quote short passages in commentary, reporting, or discussion, provided the author and the site above are credited.

The following require prior written consent: republishing this work in whole or in substantial part under another name, masthead, or brand; excerpting or modifying it in a way that alters the analysis or removes the evidence classifications; incorporating any portion of it into a paid, commissioned, or client-facing deliverable; resale or distribution behind a paywall; and use of this document as training data for machine learning models.

Commissioned Market Assessment reports, commercial licensing, and briefings on this material are available.

This document is analysis, not investment, legal, or procurement advice. It was not commissioned, reviewed, or funded by any company named in it, and the author holds no position in any of them.

Market Assessment reports, commissioned research, and briefings: **davidsoden.com/market-assessment**