





SBOM for AI (AIBOM) Tiger Team working group

Logistics

Time: Bi-weekly, Mondays at 2 pm ET | **Location:** Zoom | **Minutes:** This doc below ([link for sharing](#))

<p>Meeting Details</p> <p>Join Zoom Meeting https://zoom.us/j/99364733823?pwd=4aC6gbl1LppjBP3AuaSzr1XNHky4YP.1</p> <p>Meeting ID: 993 6473 3823 Passcode: 283563</p> <p>Bi-weekly, Mondays at 02:00 PM Eastern Time</p> <ul style="list-style-type: none">- Every 2 weeks, starting Sept 16th, 2024- Use Cases stand-up meeting is weekly, Wed at noon ET- Google Calendar <p>Please download and import the iCalendar (.ics) files to your calendar system: click here</p>	<p>AIBOM TT Leads</p> <ul style="list-style-type: none">• Helen Oakley• Daniel Bardenstein• Dmitry Raidman <p>Communication</p> <ul style="list-style-type: none">• Join Google Group: https://groups.google.com/g/aibom-tt• Join Slack: bomworkingcommunity.slack.com<ul style="list-style-type: none">◦ https://join.slack.com/t/bomworkingcommunity/shared_invite/zt-28lqtanxg~~QDeWQeZ1oPvHv34_FfEaA• AIBOM GitHub (org and all repos): https://github.com/aibom-squad• Shared google drive:  AIBOM TT docs
<p>Guiding Principles and Overview</p> <p>GitHub repo for this working group: https://github.com/aibom-squad/SBOM-for-AI-Tiger-Team</p> <ul style="list-style-type: none">- Overview of the group & guiding principles- Key activities and additional resources (SBOM for AI examples and other artifacts)	<p>Use Cases - working documents folder</p> <p> AIBOM Use Cases working docs</p> <p>Overall Use Cases - for publishing (Google doc)</p> <ul style="list-style-type: none">-  SBOM for AI Use Cases (Final Draft v0.3).docx- Consolidated doc of use cases for publishing <p>AIBOM Fields (work docs folder)</p> <p> AIBOM Fields working docs</p>

Minutes

(last meeting at the top)

Meeting Date: : 29 Sept 2025

Attendees

- Helen Oakley
- Elyas Rashno
- Raymond Sheh
- Anusha
- Brin Curcaneanu
- Dmitry Raidman
- Enoch Wang
- Michael
- Victor Lu
- Allan Friedman
- Yotam Perkal

Notes

- Presentation by SPDX colleagues, Gopi and Elyas, on SPDX field mapping to use cases
 - Based on this doc: [+ AI&Data_Mapping](#)
 - SPDX paper is planned for submission today and will be shared with the community

Meeting Date: : 18 Aug 2025

Attendees

- (add your name here)
- Helen Oakley
- Raymond Sheh
- Victoria Ontiveros
- Jono Spring
- Elyas Rashno
- Jacob Friedman
- Anthony Stephens
- John Nuckles ODNI
- Divjot Bawa
- Dmitry Raidman

Notes

- Publishing use cases under CISA lib

- Undefined timeline; awaiting for new leadership
- Next steps for fields mapping
 - Proposed to have a workstream under OWASP GenAI Security Project: <https://genai.owasp.org/>
 - Generally accepted by the Tiger Team
 - Concern: not everyone is on slack
 - Fields review: focus on “minimum fields” - the fields that *serve* the use cases (and then request those to add to CDX/SPDX if not available)
 - Elyas’s research doc: [AI&Data_Mapping](#)
 - Divjot’s mapping doc: [SBOMforAIMinimumElementsCrosswalkExercise_7.7.xlsx](#)
 - Discussion: what kind of information organizations are looking to obtain with SBOM for AI (based on what we know today)?
 - Dmitry: model parameters, size
 - Jacob: base model, architecture

Next Steps

- Helen to sync with OWASP GenAI leadership and will provide an update at next call in this round
- Dmitry, Helen and Jacob to meet to brainstorm base model and architecture data for AI SBOMs

Meeting Date: August 4, 2025

Attendees

- Helen Oakley
- Anthony Stephens
- Marec Grac
- Divjot Bawa
- Victor Lu
- John Nuckles ODNI

Notes

- Briefly discussed last call and decides to end the call today and regroup on two weeks


Meeting Date: 21 Jul 2025

Attendees

- John Nuckles ODNI
- Helen Oakley

- Allan Friedman (CISA)
- Christine Lai (CISA)
- Brindusa Curaneanu
- Elyas Rashno (Queen's University)
- John Cavanaugh (IISC)
- Marek Grac (Red Hat)
- Yotam Perkal (Zscaler)
- Bob Martin (MITRE)
- Raymond Sheh
- Dmitry Raidman

Agenda

- Update on Use Cases doc publication:
 - <https://github.com/aibom-squad/SBOM-for-AI-Tiger-Team#sbom-for-ai-use-cases>
 - Helen to add PDF version on GitHub
 - Shall we have 1-2 sentences on what SBOM for AI is? A short “go-to” definition.
 - Allan:
 - https://www.bsi.bund.de/SharedDocs/Downloads/EN/BSI/KI/SBOM-for-AI_Food-for-thoughts.pdf
- Fields mapping
 - **Project goal:** Publish “Minimum Fields” for SBOM for AI
 - Deliverable: top 10 fields are defined in a document (across use cases)
 - Question: dependencies vs. knowledge graph
 - What is the scope of these 10 fields (SBOM standard fields included vs. AI-specific fields only)
 - Reference from SPDX:
 - https://www.linuxfoundation.org/research/ai-bom?hss_channel=lcp-208777
 - <https://spec.c2pa.org/specifications/specifications/2.2/index.html> <https://cawg.io/>
 - <https://mlcommons.org/working-groups/data/croissant/>
 - <https://aibom.org/schemas/>
 - <https://www.w3.org/TR/vc-data-model-2.0/>
 - Discussion regarding the “cost” of the model (resource consumption, parameters)
 - **Approach:** to be discussed
 - Example: map fields for use cases, then find common fields
 - Elyas Rashno is proposed to share a document with use cases from other research paper (there is overlap with these use cases) and field mapping for AI SBOM
 - Divjot's mapping:  SBOMforAIMinimumElementsCrosswalkExercise_7.7.xlsx
- Industry updates:
 - BlackHat & DEFCON 33 (Las Vegas)

- SBOM meetup on Aug.5 at 12-2pm
- AI SBOM Generation demo at DEFCON AppSec Village Arsenal Aug.8
- Vulnerabilities: <https://research.eye.security/sharepoint-under-siege/>

Next Steps

- Elyas and Helen to show field mapping based on other research papers and current use cases we have
 - [+ AI&Data_Mapping](#)
- Allan & Divjot: to prepare 1-2 sentences on what SBOM for AI is
- Helen to add PDF doc to GitHub (done - [here](#))
- Next meeting is on Aug.4th (even if some are away, we keep it as a working session)

Meeting Date: 7 Jul 2025

Attendees

- Allan Friedman
- Daniel Bardenstein
- Raymond Sheh
- Yotam Perkal
- Dmitry Raidman
- Elyas Rashno
- Divjot Bawa
- Jyoti Wadhwa

Agenda

- **Guest speaker:** Gareth Richards from UK's TIABOM efforts
- Update on the use cases doc process: see last meeting minutes below
- Field mapping: see last meeting minutes below
- Industry updates

Next Steps

- ...



Meeting Date: 23 Jun 2025

Attendees

- Helen Oakley
- Raymond Sheh
- Courtenay Ngo(Microsoft)
- Elyas Rashno (Queen's University)
- Bob Martin (MITRE)

- John Nuckles (ODNI)
- Dmitry Raidman
- Marek Grac
- Brian Hierholzer
- Allan Friedman

Agenda

- Update on the use cases doc process
 - **June 23rd** – SBOM for AI TT call & Final draft announcement
 - GitHub pull request with an update is submitted; would request for a minor follow-up request once contributors have submitted their names:
<https://github.com/SBOM-Community/documents>
- Field mapping:
 - Helen showed spreadsheet with fields as an idea how to approach the mapping and field prioritization; the ask to the community is to review and come back with the feedback on best ways to proceed; current starters (we can build on those, or decide to start fresh with some other approach):
 -  AIBOM fields mapping.xlsx
 -  Proposed AIBOM TT Scope
 - Approach: proposal to start with use case by use case (ideal, authors of those use cases would be involved) to identify what kind of metadata could be required for those use cases; then we could map that metadata to existing (of future) fields within technical format of SBOM for AI; based on mapped fields across use cases, we can identify the most important AI-relevant fields (which could be then released as an extension of current SBOM framing document or a separate one for SBOM for AI)
 - Brian H. proposed to manage field mapping on GitHub .md to potentially help with the modeling of data and relationships between documents through a tool (in development)
 - Discussion on modal card data quality and challenges: ideally we should have standardized model card format and metadata, with automated generation (perhaps, based on SBOM for AI output)
 - <https://arxiv.org/abs/2402.05160>
- [TAIBOM](#) (UK effort for AIBOM): Allan proposed to invite them over for collaboration

Next Steps

- Announcement of the doc: Allan will follow-up on the steps; and will also follow-up on GitHub pull requests
- Field mapping: Helen/Dmitry/Daniel to define working items for addressing fields mapping exercise

- Ask to the community: to review current “starter” docs and come back with feedback and ideas on how shall we proceed with the field mapping work

Meeting Date: 9 Jun 2025

Attendees

- (add your name here)
- Allan Friedman
- Helen Oakley
- Daniel Bardenstein
- Marek Grac
- Raymond Sheh
- Animesh Pattanayak
- Courtenay Ngo
- Yotam Perkal
- Elyas Rashno
- Bob Martin

Agenda

- Update on the use cases doc process
 - **June 9th** – Meet with AI TT to review community feedback and adopt/reject, continuing with the document update
 - Draft for community review:
 - [w SBOM for AI Use Cases \(Community Comment\) .docx](#)
 - Contributors / reviews, please add your name here: [+ Contributions](#)
 - Feedback form: <https://forms.gle/CHrEsAQSzrE1Kkx96>
 - Results spreadsheet:
 - [+ SBOM for AI Use Cases - Community Feedback \(Responses\)](#)
- Field mapping (will discuss during next call)
- Industry updates
 - Potentially, a virtual SBOM-a-Rama might be organized by CISA this summer

Next Steps

- **June 16th** – Tentative publication on Github page; final touches/fixes during the following week
 - <https://github.com/SBOM-Community/documents> → add a link to read-only version on this page for now

- **June 23rd** – AI SBOM TT call & Final draft announcement

Meeting Date: 26 May 2025

Attendees

- (add your name here)
- Is there a meeting now?

Agenda

- Update on the use cases doc process
 - Draft for community review:
 - [w SBOM for AI Use Cases \(Community Comment\) .docx](#)
 - Feedback form: <https://forms.gle/CHrEsAQS zrE1Kkx96>
- tbd

Next Steps

- ...

Meeting Date: 12 May 2025

Attendees

- Helen Oakley
- Animesh Pattanayak (PNNL)
- Raymond Sheh
- Bob Martin (MITRE)
- Elyas Rashno
- Sumit Giri
- John Nuckles (ODNI)
- Shafia Zubair
- Allan Friedman

Agenda

- **Update from RSAC workshop:**
 - <https://github.com/aibom-squad/AIBOM-RSA-2025>
 - Slides and recording are to be shared (links will be updated on github)
- **Use Case doc update:** [w AI SBOM Use Cases \(Community Comment\) .docx](#)

- Document release plan:
 - **May 12th** – Use AI SBOM TT meeting to update on comment consolidation; finish any lingering conversations asynchronously over email, if necessary, over the course of the next week
 - **May 16th** – **Helen** to email to the AI SBOM TT with the update on cleaned up doc
 - **May 21st** – **Helen** and/or **Allan** send out a note to AI SBOM TT indicating that all edits have been adjudicated, and the document should be shared externally for wider community feedback
 - **May 21st** – Start of public comments: everyone shares on LinkedIn, etc.
 - **June 4th** – Close period for external comments; AI SBOM TT volunteers to adjudicate and begin consolidating comments the first week of June
 - **June 9th** – Meet with AI TT to review community feedback and adopt/reject, continuing with the document update
 - **June 16th** – Tentative publication on Github page; final touches/fixes during the following week
 - **June 23rd** – AI SBOM TT call & Final draft announcement
- Doc usability/readability update:
 - Roles/personas: **Helen** is updating the doc
 - Summary/problem statement to be consolidated/cleaned-up
- SBOM community comments: LinkedIn / public review
 - How to collect feedback? Options:
 - **Google Forms** (preferred; pros: can structure and group feedback)
 - **Google Forms option is selected**; Helen will create first draft
 - Google Doc comments (cons: too difficult to track)
 - Github pull requests (cons: might not be straight-forward for non-developers)
 - **Document to be published as final draft on June 23rd**
 - Published on: <https://github.com/SBOM-Community>
 - Create new repo and link in “documents” repo too
- **Naming: “AI SBOM” vs “SBOM for AI”**
 - Allan: there will be an email regarding naming decision
- **Fields mapping** - open discussion for following calls:
 - Approach is to scope the next stage based on following questions:
 - Which fields provide the most value for risk identification?
 - How do they connect to specific use cases?
- **Industry updates**
 - Try [AI SBOM Generator](#) (in CycloneDX format for models on Hugging Face)
 - Read blog about it [here](#)

Next Steps

- **Helen** to drive doc cleanup until May 16th and hand it over to Daniel & Dmitry (Allan will support)
- **Everyone**—starting May 21st: spread the word across your network about broad community review of the document (May 21 - June 4)

Meeting Date: 28 Apr 2025

Meeting on April 28th will NOT take place due to RSA.

- RSA / AIBOM workshop at RSA 2025—register here: <https://lu.ma/d9sry4mi>

Meeting Date: 14 Apr 2025

Attendees

- Daniel Bardenstein
- Raymond Sheh
- Scott Bowman (INL)
- Elyas Rashno (Queen's University)
- Allan Friedman
- Divjot Bawa
- Sandeep Purewal (HiddenLayer)
- Arthit Suriyawongkul (ADAPT Centre, Trinity College Dublin)
- Nick Mistry (Lineaje)
- Kim Sevenz

Agenda

- VulnCon
- Use Case document
 - **April 18th:** end of CISA SBOM comment period
 - **Helen + friends:** learn & integrate comment adjudication
 - Divjot Bawa
 - Raymond Sheh
 - Arthit Suriyawongkul (Use Case: Compliance)
 - Elyas Rashno
 - **April 25th:** comment adjudication
 - **Public comment period:** April 28th - May 16th.

- **PC adjudication:** May 19-30th
- **Final release/publication:** ~June 2nd
- RSA / AIBOM workshop at RSA 2025—register here: <https://lu.ma/d9sry4mi>
 - One SBOM-related talk on healthcare
- Call for Paper: Nordic Software Security Summit – Deadline: 1 May
<https://nsss.se/y2025/call-for-papers/>
- AI SBOM Fields
 - **Objective 1:** Identify a baseline set of fields for AI-related components in SBOMs
 - “Minimum elements”: is this different from traditional minimum elements
 - “Baseline elements”: what is good / sufficiently robust in terms of comprehensiveness of AI-related components in an SBOM
 - Baseline Element <-> SPDX mapping table (Proposal - Not Official):
https://docs.google.com/document/d/1H9p9ipdlhX2ErLi5Y6zhbmGbalA1ChYn_kgIS4Eq88o/edit?tab=t.0
 - CycloneDX (including the ML BOM) <https://cyclonedx.org/docs/1.6/json/>
 - For example,
https://cyclonedx.org/docs/1.6/json/#components_items_modelCard
 - SPDX AI Profile <https://spdx.github.io/spdx-spec/v3.0.1/model/AI/AI/>
 - How to use SPDX AI and Dataset Profiles
<https://www.linuxfoundation.org/research/ai-bom>

Next Steps

- **Helen + friends:** work on setting up comment adjudication process
- **Daniel + Arthit + Nick + Kim + Divjot + Raymond + Brin + Sandeep:** Comment Adjudication
 - Some tasks can be done in parallel (obvious typos and grammars can be done without the wait for the adjudication process to be settled)
- First draft of elements:

Meeting Date: 31 Mar 2025

Attendees

- Daniel Bardenstein (Manifest)
- John Nuckles (ODNI)
- Allan Friedman (CISA)
- Cassie Crossley (Schneider Electric)
- Bob Martin (MITRE)
- Divjot Bawa (CISA)
- Arthit Suriyawongkul (ADAPT Centre, Trinity College Dublin)

- Emily Fox (Red Hat)
- Kim Sevenz (TCVS)
- Raymond Sheh
- Elyas Rashno (Queen's University)
- Nick Mistry (Lineaje)
-

Agenda

- **Use Cases Doc**
 - Ready for CISA SBOM Community Review
 - CISA: best way to share is drop in the Google Group with a timeline and explicit scope / types of feedback
 - A priori structure / process for comment adjudication
- RSA workshop update
 - Register on Luma: <https://lu.ma/d9sry4mi>
- **AI SBOM Fields**
 - Soft project launch today
 - Future sessions:
 - Define outcome / deliverable
 - Define scope
 - Timelines, etc.
- Other updates
-

Next Steps

- Send out Use Cases draft to CISA SBOM Community Google Group (Daniel)
- Set up time with Melissa and or Anita for comment adjudication process (Daniel / Leads)
- Create (or find existing) AI SBOM field tracker Google Sheet (Daniel)
- Mention AI SBOM workshop at next week's SBOM Community Call (Leads)

Meeting Date: 17 Mar 2025

Attendees

- Helen Oakley
- Raymond Sheh
- Jim Light - IIS
- Kim Sevenz
- Yakov Neshcheretnyy(MITRE)
- Shafia Zubair

Agenda

- Update: SBOM overall use case doc review - finalizing the updates and will send for feedback to the broader SBOM community this week
 - We will keep the bi-weekly calls to review and follow-up on community feedback
- RSA AIBOM workshop:
 - Date: April 28th afternoon (12 - 3pm) in-person during RSA, hybrid - tbd
 - Scope: tbd, current proposal is for
 - Update on Use Case
 - Field mapping to use cases

Next Steps

- SBOM overall Use Cases doc feedback from broader community
- RSA AIBOM workshop preparation (Helen, Daniel, Dmitry)

Meeting Date: 3 Mar 2025

Attendees

- Helen Oakley
- Daniel Bardenstein
- Enoch Wang (LLNL)
- Lance McCaleb (DoD)
- William Malik - IIS
- Raymond Sheh
- John Nuckles (ODNI)
- John Cavanaugh - IIS
- Emily Fox (Red Hat)
- Divjot Bawa (CISA)
- Dmitry Raidman

Agenda

- Consolidated document: [AI SBOM Use Cases \(overall\).docx](#)
 - Includes not only all selected use cases but also an intro (exec summary, overview) and conclusion (overall recommendations)
 - Reviewers to comment in the consolidated document (not in individual working docs)

- (Review CNCF example -> Emily Fox)
<https://github.com/cncf/toc/blob/main/tags/resources/toc-supporting-guides/tech-papers.md>
 - Specifically the adjudication section - happy to help guide and navigate adjudication as well since I've done this quite a few times.
 - Other notes: Administrative comments should be able to be resolved on the fly as they come in (spelling, grammar, clarity without impacting intent or content) - this reduces the adjudication burden if managed as they occur. Substantive comments require more work by the authors but fundamentally don't alter the content, rather it may be reorganizing material or adding new content that was overlooked. These also call out assumptions that need to be addressed. Critical comments require coming back to the group to address and will result in changes to the intent and content. In some cases it may be presenting dual perspectives for the reader to come to their own conclusion.
- AI SBOM fields
 - Proper kickoff next AI SBOM Tiger Team
- Conferences
 - VulnCon
 - Attendees: Daniel, Dmitry, John C
 - RSA
 - Determine who from the TT will be at RSA
 - TT Leads to put together proposal for AI SBOM field workshop at RSA

Next Steps:

- Final review by Thursday March 6th (AI SBOM tiger team).
 - Present at next CISA Community meeting on March 10th
 - Determine adjudication process & history
 - Helen to send email to AI SBOM Tiger Team (done on March 3rd)
- TT Leads to put together proposal for AI SBOM field workshop at RSA

Meeting Date: 18 Feb 2025

Attendees

- Daniel Bardenstein
- Helen Oakley
- Raymond Sheh
- Dmitry Raidman
- Allan Friedman

- Divjot Bawa
- Cassie Crossley

Agenda

- S4 update: <https://s4xevents.com/>
 - Helen's talk: went well + recorded!
 - Required some education around SBOMs
 - Cassie: interest around AI SBOM from Dept of Energy, energy/infra sector
 - Helen's smaller talk on AI minimum field at SAP
- Use Cases Update
 - We will consolidate the use cases into a master doc and then review for consistency, etc.
- AIBOM fields
 - Schedule Kickoff // Set scope for future session
- RSA AIBOM workshop plan
- Update from CISA
- Policy updates
 - Guide: OWASP released Agentic AI Threats & Mitigations Guide
 - Question: how does BOM apply to AI Agents, how does it differ?
 - Different types of components in BOMs – type = "AI Agent" ?
 -

Action Items

- Consolidate all UCs into a single doc (Helen)
- Review UC 1 and 6 (Daniel)
- Review UC 5 (Secure by Design) - all
- Review UC X: Ray to de-scope

Meeting Date: 3 Feb 2025

Attendees

- Daniel Bardenstein
- Raymond Sheh
- Allan Friedman
- Helen Oakley
- Emily Fox
- Cassie Crossley
- Ian Dunbar-Hall

- Nick Mistry
- John Cavanaugh - IIS
- Kim Sevenz
- Kenneth Peeples (Red Hat)
- Dmitry Raidman
- John Nuckles (ODNI)
- Bob Martin (MITRE)

Attendees

- Use Cases Update:
 - Use case status doc: [AIBOM Use Cases - working document](#)
 - SBD: update from Shafia - to be completed by EOD
 - UC7 ready for review
 - UCX offered for consideration
 - Drop IP Use Case -> put into Compliance Use Case?
 - Allan: is it such a dominant issue that we should pull out?
 - Daniel to see if he can fill out IP, if not - then we may need to de-scope it
 - Allan: Not sure if there's existing SBOM docs/use cases for licensing
 - UC5 proposed to scope out, "secure by design".
 - Roundtable:
 - UC 1-3 ready for review
 - UC 4: ready for internal review (readability)
 - UC 5: SBD (wait for Shafia) - may need to change the title
 - UC 6: IP / Fair usage. Daniel to work on
 - UC 7/8: Reformatted. Ready for community review
 - Perhaps read through on UC 8 inclusion
 - UC X: Raymond's idea, ready for review
 - https://docs.google.com/document/d/1vi778T7gK_kac-DIac06mdJ1z8ngy4c-f8jQud3nnfo/edit?tab=t.0#heading=h.2jb0d4kbc35
 - Glossary: what other terms to add?
- Ping Allan/CISA for how to integrate AI into other CISA guidance?
 - Framing doc
 - Cisa.gov
 - SBD
 - Other AI docs
- **AIBOM fields:**
 - Folder with working docs: [AIBOM Fields working docs](#)
 - To continue discussion during the next meeting.
- **Policy updates**
 - [USPTO AI Strategy](#)
 - [NCSC CoP](#)

- **General updates**
 - Re-naming "AIBOM": "AI SBOM" (13) vs "SBOM for AI"(4)
 - Voting: preference is towards the "AI SBOM" as a result of the vote during the call

Next Steps:

- **Everyone:** Look at Use Cases 1-4, 7, X, and glossary "Formatted" with comments and suggested changes before the next meeting. (rows with a **green Ready for Review**)
 - Status and doc links are here: [AIBOM Use Cases - working document](#)
- **Daniel:** US6 (IP & Fair use) review and see if can fill in the gaps
- **Shafia:** US5 (Secure design) to notify on completion; propose another title
- **Daniel/Helen/Dmitry:** Ping Allan & CISA folks in a few weeks re updating CISA BOM documentation
- **Daniel/Helen/Dmitry:** Move meeting to be moved to Feb 18th (from Monday, Feb.17)

Meeting Date: 21 Jan 2025

Attendees (please add yourself)

- Helen Oakley
- Raymond Sheh
- Nick Mistry
- Jim light
- William Malik
- Jonathan Comella (CISA)
- Barbara Shelton (CISA)
- Zach Hambrice
- John Nuckles (ODNI)
- Cassie Crossley (Schneider Electric)
- Yotam Perkal
- John Cavanaugh - IIS

Agenda

- Review and feedback of proposed one-pager for "SBOM for AI":
 - [SBOM for AI: Introduction](#)
 - Notes:
 - Perhaps, "SBOM for AI" is not the best term

- The proposal is to ask CISA to provide a common description for an SBOM and that it includes all other variations like AIBOM, HBOM, etc. - in the framing document:
<https://www.cisa.gov/sites/default/files/2024-10/SBOM%20Framing%20Software%20Component%20Transparency%202024.pdf>
- Use Cases updates
 - Standup for writers is on Wed's
 - Community review: Victoria / Barbara (CISA) would send the notification for community review
- Industry updates
 - New EO 14144:
<https://www.federalregister.gov/documents/2025/01/17/2025-01470/strengthening-and-promoting-innovation-in-the-nations-cybersecurity>
 - see Section 10 Definitions: section 10 c the term "artificial intelligence" or AI has the meaning set forth in 15 U.S.C. 9401(3)

Next Steps

- Helen: reach out to CISA regarding updating the framing document for including SBOM scope (eg., AIBOM, etc.)
- Helen: email writers to confirm readiness for community review

Meeting Date: 6 Jan 2025

Attendees (please add yourself)

- Daniel Bardenstein
- Helen Oakley
- Emily Fox (note taker extraordinaire)
- Bob Martin
-
- SA)

Agenda

- General updates
 - **NEXT CALL is moved to Jan.21** (Tue) due to public holiday on US on Jan.20th (Helen to send the email to the group)
 - Re-cap on last meeting and follow-up actions

- Standup calls on Wednesdays at 12-12:30pm ET (download from the same calendar - link above)
- Have you joined our Slack? It's a good way to stay in-touch between the calls.
- Industry updates (policies, events)
 - AIBOM talk at <https://s4xevents.com/agenda/>
 - ...
- Use case discussion (Daniel)
 - Status review
 - Timeline / next steps

AIBOM fields

-

Next Steps

... **(Helen)**

First Draft of AIBOM vs SBOM for AI, what's in it, etc. **(Helen)**

-

Meeting Date: 9 Dec 2024

Attendees (please add yourself)

- Helen Oakley

-

Agenda

- General updates
 - Standup calls:
 -
- Industry updates
- Use case discussion

Next Steps

- ...
- Helen to send a reminder through google group
- Dmitry will setup a call for use case writers for the week of New Years (before Jan.6 deadline)

Meeting Date: 25 Nov 2024

Attendees (please add yourself),

- ...
- Helen Oakley

Agenda

- General updates:
 - Break-outs calls
 - Standup for use case writers
 - AIBOM Fields workstream
 - Timelines: call for feedback from use case writers on ETA
 - Please add your ETA to column F in this document:
https://docs.google.com/spreadsheets/d/1165T3oQvxOjKvcSbR-txmKAluXffRZ_75yZAJyeACxY/edit?usp=sharing
 - **Community agreement for deadline for first draft: January 6th, 2025 (AIBOM TT WS call)**
 - Add yourself to the Google Group for AIBOM communication and calendar invites:
<https://groups.google.com/g/aibom-tt>
- Industry updates
 - Policies, papers, events
 - SPDX publication for AIBOM: <https://spdx.dev/implementing-an-ai-bom/>
- Use case discussion
 - Community round table

Next Steps

- Standup calls for use cases (Dmitry will email invite)
- Use case writers: attend standup calls and aim for Jan.6 deadline for your first draft
- **Need volunteers to help with use case #6 "Intellectual Property & Fair usage"!**

Meeting Date: 12 Nov 2024

Attendees (please add yourself)

- Helen Oakley
-

- Dmitry Raidman

Agenda

- Follow-up from last meeting
 - Potential deadline for draft readiness? (Daniel): **Feb 3 release.**
 - Identify a reference UC document (Raymond + DJ)
 - **<https://www.cisa.gov/sites/default/files/2024-05/SBOM%20Sharing%20Primer.pdf>**
 - TT email list (Daniel/Helen): <https://groups.google.com/g/aibom-tt>
 - Please ask to join this group to add yourself.
- Discussion:
 - Google group (Helen)
 - Will figure out access
 - Back to weekly meetings
- AIBOM fields
 - Project lead(s): Dmitry, Daniel, Ronan, Karen B, Raymond
 - **Proposed AIBOM TT Scope**
- Updates
 - Policy
 - Election impact
 - [Texas Draft bill](#)
 - Cassie: assume that "AI Transparency" in regulation/policies implies AIBOMs
 - Only 7 global policies explicitly mention (AI)BOMs

Next Steps

- UC milestone delivery plan/timeline: (Daniel)
- Send out weekly vs bi-weekly poll (Dmitry)
- Schedule initial AIBOM Fields work (Dmitry)
- Google group / list of people and contacts (Helen)
 - **Fixed visibility of this group** - please ask to join to add yourself:
<https://groups.google.com/g/aibom-tt>

Meeting Date: 28 Oct 2024

Attendees (please add yourself)

- Daniel Bardenstein
- Helen Oakley

- Laura Randall

Agenda

- Quick catch-up from Helen:
 - Overview from AIBOM Workshop
 - Deeper dive into fields <> proposal mapping
 - **Field mapping may not be top priority right now, but people are free to contribute**
 - **Google Drive - [link](#)**
- Status update: Use Cases, next steps?
 - Helen has re-structured UCs <> Google Docs, ownership, etc.
 - [Updated tracker](#)
 - Publishing:
 - Would take a while for publishing within CISA framework. Would be helpful to prep draft for CISA review
 - Victoria: how long do we want to allow community review for?
 - Raymond: template of what a good UC Document looks like, style-wise?
 - Karen: perhaps examples one of Standards Compliance
 - **Next steps:**
 - Can UC owners provide estimates for community review? Or provide deadlines?
 - Duration of community review: 2 weeks?
 - **Potential deadline for draft readiness? (Daniel)**
 - **Identify a reference UC document (TODO)**
 - Victoria:
 - https://www.ntia.gov/sites/default/files/publications/ntia_sbom_use_cases_roles_benefits-nov2019_0.pdf or
 - <https://www.cisa.gov/resources-tools/resources/sbom-sharing-primer> or
 - <https://www.cisa.gov/resources-tools/resources/vulnerability-explorability-exchange-vex-use-case-document-april-2022>
 - **TT email list for outreach**
- Status update: AIBOM fields, next steps?
 - https://docs.google.com/document/d/16UO9gOfVBj_JRCgo01aXwt8A0cVTaQEA_NywB7PJsrGI/edit?tab=t.0#heading=h.2o7mgm6qoqcu
 - Perhaps stating the AIBOMs are inclusive of "Model Card" information, so not limited in terms of where storage
 - E.g. OCI storage / associated with containers
 - Have discussed model card generation with/from AIBOM
 - SPDX talking with Hugging Face; at minimum, link to model card in AIBOM
 - Github working on something similar

- https://github.com/huggingface/huggingface_hub/blob/main/src/huggingface_hub/templates/modelcard_template.md
- **Follow Up with Karen re Hugging Face contacts.**
- Reg updates:
 - [New AI Executive Order in US](#)
 - [OMB Guidance on AI Procurement for USG](#)
 - [NVIDIA announcement](#)
- General updates:
 - Karen Bennet: new EO on synthetic data. Perhaps as use case

Next Steps

- Potential deadline for draft readiness? (Daniel)
- Identify a reference UC document (Raymond + DJ)
- TT email list (Daniel/Helen)
- Keep working on use case documents

NOTE: Meeting on October 14th is skipped due to a federal holiday in US and public holiday in Canada

Meeting Date: 30 Sept 2024

Attendees (please add yourself)

- Amy Villaseñor
-
- Dmitry Raidman - Cybeats
- Helen Oakley - SAP

Agenda

- AIBOM minimum set of fields (eg., for mode card) vs. MVP (for Compliance & Vuln. Mgmt use cases)
 - [Using the Proposed AIBOM TT Scope doc to document the AIBOM fields mapping table](#)
 - [The table now has some examples for each field](#)
 - [AIBOM MVP will include baseline fields. However, the MVP will include more fields since it will focus on the “Compliance” and “Incident Response” use cases](#)
 - [How do the current CycloneDX ml-bom fields map to the AIBOM or what is the difference?](#)
- Use cases timeline, and milestones
 - [Every use case team should set their own timeline so that the rest of the TT can review them. Team members to add their timeline to the spreadsheet](#)

- Please add your names on the agenda once you are ready to present. Make sure to note this in the spreadsheet
- Use case working docs folder on AIBOM TT drive
 - AIBOM Use Cases working docs folder: <https://drive.google.com/drive/folders/1-X1YPCtLJRtSXyVfDAOuCbQTsYFsdKkF?usp=sharing>
 - Individual use cases are linked in the spreadsheet and the overall doc
 - Overall AIBOM TT docs: <https://drive.google.com/drive/folders/11NmamWuNmyuudaW1LBvu6raUb63xJ4L2?usp=sharing>
 - Each use case will have its own google document which will then be consolidated into the main document (see links above).
 - For now we will change the main doc to be read only to make it clear that writing should take place in individual documents
 - For next call, @Victor Lu to present use cases 7&8 so the TT can decide if they should be consolidated or remain separate
- Regulations update: SB 1047 veto, but passing several other bills for AI. Read more: <https://www.gov.ca.gov/2024/09/29/governor-newsom-announces-new-initiatives-to-advance-safe-and-responsible-ai-protect-californians/>

Table below migrated to [Proposed AIBOM TT Scope](#) please add directly to linked document

[Deprecated] Meeting Date: 16 Sept 2024

Attendees (please add yourself)

- Amy Villaseñor
-
- Dmitry Raidman
- Helen Oakley
- Cassie Crossley
-

Agenda

To organizers: please record this call as we have a guest speaker.

- **Guest speaker:** Allan Friedman with an overview for documentation publishing
- **Update from AIBOM workshop at SBOM-a-Rama** - see highlights on GitHub, with linked slides and doc from the workshop: https://github.com/aibom-squad/SBOM-a-Rama_AIBOM_Fall2024.

- The recording will be published later this week, link will be added to GitHub too (or subscribe to this channel to watch:
<http://www.youtube.com/@SoftwareSupplyChainSecurity>)
- Workshop Recording: <https://youtu.be/CadtocSfX6E?feature=shared>
- Attendance for live event: there were ~30 people in the room, and ~40 people online (some people in the room connected online too)
- **Use case progress:** Document structure, updates from owners, open discussion
 - Suggestions from SBOM-a-Rama session: what does the use case relate to?

Notes

- Discussion...
- Action items...

Meeting Date: Aug 26, 2024

- Attendees (please add yourself)
-
- [Amy Villaseñor](#)
- Dmitry Raidman (Cybeats)
- Helen Oakley

Agenda

- Intro & calling for a volunteer to take notes
- AIBOM TT meeting series:
 - Current series will be canceled after this call
 - Next AIBOM call will be at SBOM-a-Rama on Sept. 12th (remote access, or in-person for the AIBOM workshop)
 - New series will be scheduled bi-weekly, starting Sept. 16th (new calendar invite and zoom details will be provided and included on this document)
- Agenda for AIBOM Workshop at SBOM-a-Rama:
 - New GitHub page:
https://github.com/aibom-squad/SBOM-a-Rama_AIBOM_Fall2024
 - Registration for AIBOM workshop (virtual/in-person): <https://lu.ma/ic56tefk>
 - Helen to add goal to the GH link: goal is to define next steps once we finish use case development
- Community updates
 - Policies, news, events (round table from the community)
 - Resource shared: <https://airisk.mit.edu/>
- Use Case discussion:

- Breaking down use case (current) doc into separate docs for better structure and usability of working:
 - Option: break down for the working on the use case, and then we can consolidate into a concise doc
 - Helen will create docs and will move content, and will provide links in excel
 - Examples for use cases:
 - <https://www.cisa.gov/sites/default/files/2024-05/Software%20Transparency%20in%20SaaS%20Environments.pdf>
- Contributors and reviewers, add your name into the document at the bottom, including your organization
- Follow-up on comments in the word doc for Use Cases, any other questions / concerns from others on current content
 - eg. EU definitions from AI Act (<https://artificialintelligenceact.eu/article/3/>) and how those are mapping into SBOM definitions (eg., here: https://www.ntia.gov/sites/default/files/publications/ntia_sbom_use_cases_roles_benefits-nov2019_0.pdf)
 - Should we have Terms and Definitions mapping/description at the beginning of AIBOM Use Case doc?
- Review the objectives the volunteers filled out
 - *Anyone volunteers to lead their use case discussion?*
 - The team decided to

Notes

- Discussion
 - ...
- Action Items
 - ...

Meeting Date: Aug 12, 2024

Attendees (please add yourself)

- (add yourself here)
- Helen Oakley
- John Cavanaugh - IIS
- Amy Villasenor

- Ixchel Ruiz (Karakun)
- Shamik Chaudhuri - Google
- Raymond Sheh
- Allan Friedman (CISA)
- Nicholas Vidovich (Finite State)
- Alex Sharpe
- Georgianna Shea
- Dmitry Raidman
- Peter Armstrong
- Victor Lu
- Ricardo Reyes

Notes

- Agenda (see slides)
 - Announcements
 - AIBOM workshop at SBOM-a-rama
 - Registration link: <https://lu.ma/ic56tefk>
 - Limited in person spots, will also have a virtual component
 - Use cases agenda:
 - Discuss the “Compliance” use case
- Ameya Naik is leading the discussion for this call
- Reminder on guiding principles (outcome-driven discussion, etc.)
- Discussion of Compliance use case in the Word document ([here](#))
 - Objective:
 - Proposed initial objective text (see Word doc)
 - Question: what AI regulations are covering supply chain aspects?
 - Currently, AI regulations are not explicitly covering supply chain, however, things like accuracy of AI (the transparency of how it’s accurate) would be covered by AIBOM. Also, traditional software supply chain security requirements would also apply to AI systems.
 - Reference for EU AI Act: https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=OJ:L_202401689#d1e38-136-1
 - Transparency requirements: <https://www.europarl.europa.eu/topics/en/article/20230601STO93804/eu-ai-act-first-regulation-on-artificial-intelligence>
 - Whether the AI BOM is the place to solve for this (or parts of it) is probably an open question.
 - Generative AI, like ChatGPT, will not be classified as high-risk, but will have to comply with transparency requirements and EU copyright law:

- Disclosing that the content was generated by AI
 - Designing the model to prevent it from generating illegal content
 - Publishing summaries of copyrighted data used for training
- Scope of this use case is within compliance, but not so much of ethics (ethics and fairness are part of another use case)
 - We shall outline the scope / definition of “compliance” part
- “EU AI Act that differentiate AI Provider from AI Deployer” - Resource reference:
 - <https://github.com/bact/stav/blob/main/presentations/ai-accountability-taonomy-arthit-oss-na-20240416.pdf>
- Personas:
 - Proposed terms: AI Provider, AI Deployer, Regulator; regulatory bodies, devs and engineers, internal auditor/compliance teams, business user
 - Would 3rd party auditors be part of “regulatory bodies” or “internal audit”?
 - Likely to be part of extended scope of “regulatory auditors” (term to be defined)
 - Business user: how is it relevant to compliance use case?
 - Proposal to separate “business owner” from “data owner”
 - “Business User” is currently remanded to “Business Owner”
 - “Data Owner”: final authority of how data is collected and used; other relevant terms - data custodian and data stewards
 - To consider: datasets for training, sensitive data usage, etc.
 - Following the SBOM taxonomy can we say we have AIBOM producers and AIBOM consumers?
 - Proposal to categorize personas into the groups of “producers” and “consumers”
 - Digital Asset Owner – CISO Chief Information Security Officer or BISO Business Information Security Official. Responsible for the use of sensitive data and security of sensitive data. Digital Asset Owners may oversee data operations of an entire organizations or business entities, or oversee data operations of lines of business with P&L responsibility, or oversee data operations of functional areas such as Payroll, AR/AP, Help Desk, Sales, Supply Chain Operations, or Quality Assurance

Actions

- Everyone: Review initial use case proposal text and add comments
- Volunteers: pick next use case to fill, use excel as supporting doc (starting point and initial use case info)

Meeting Date: Aug 5, 2024

Attendees (please add yourself)

- (add yourself here)
- Emily Fox (Red Hat)
- [Amy Villaseñor](#)
- Helen Oakley
- Scott Heimann
- Mark Orsi
- Marek Grac
- Nicholas Vidovich
- Dmitry Raidman
- John Nuckles
- Alex Sharpe

Notes

- Agenda (see slides)
 - Announcements
 - AIBOM workshop at SBOM-a-rama
 - Registration link: <https://lu.ma/ic56tefk>
 - Limited in person spots, will also have a virtual component
 - Use cases agenda:
 - Prioritize consolidated use cases
 - Review template and decide on next steps
 - Bug Fix/Defect Tracing
 - Question of how this fits into the AIBOM use case
 - Open questions:
 - Are there tunable attributes of a model that should be documented in the AI BOM?
 - If we go this route, would we need to think about how to classify all these characteristics so that it is machine readable and not just qualitative?

- Is it fair to say that AIBOM is a standalone object that is used for risk management
- Maybe the word 'bug' is incorrect, updating the use case name to 'model change management specs/lifecycle'
- Example: 2 different LLMs to generate a draft focused on a specific technology. Wanted to make sure the recipient had knowledge of the LLM and the query used
- At what point is an AIBOM generated?
- How do we capture characteristics of a model that need to be factored into a risk management decision of incident response but it's not a defect, it's not a vulnerability?
 - Example: rocket engines (2 classes, liquid fuel & solid fuel)
- We may need some research to understand the categories of items that are dynamically tuned in the wild
- We will approach the use cases in priority order (priority has been added to the use cases sheet)

Notes from the meeting chat (AI assistant):

The meeting focused on discussing and prioritizing the use cases for AIBOMs. The participants discussed the importance of capturing bug fixes and defects in the AIBOMs, but there was a difference of opinion on whether it should be included as a use case. Some participants argued that bug fixes and defects should be tracked separately, while others believed that it is important to showcase how bug fixes and defects can be used for reactively updating certain packages or dependencies.

The discussion then shifted to the concept of capturing characteristics of AI models that need to be factored into risk management decisions. Examples were given, such as the tendency of a model to classify colors differently or the configuration parameters of a model. It was suggested that these characteristics should be captured in the AIBOMs to provide transparency and enable effective risk management.

There was also a discussion on the dynamic nature of AI models and how to handle them in the AIBOMs. It was acknowledged that capturing snapshots of dynamic models could be challenging, but it was suggested that a field indicating whether a model is dynamically updated could be included in the AIBOMs.

The participants agreed to focus on shippable models as a starting point and to define the scope of the use cases. It was suggested that the use cases should be documented using a provided template and reviewed in the next iteration.

The meeting concluded with a discussion on prioritizing the use cases. The participants agreed that compliance should be the first priority, followed by vulnerability and risk assessment.

Overall, the meeting was productive in discussing the use cases for AIBOMs and setting priorities for further work. The participants provided valuable insights and suggestions for capturing bug fixes, defects, and characteristics of AI models in the AIBOMs.

- The discussion revolved around a proposed use case related to bug fixes and defect tracking in AI models.
- There was a debate on whether this use case should be included in the AIBOMs or if it should be addressed separately.
- Different perspectives were shared, with some arguing that bug fixes and defects should be treated as vulnerabilities and tracked in the AIBOMs, while others believed it should be handled outside of the AIBOMs.
- The idea of capturing the history or unknown weaknesses of AI models was brought up, with the need to categorize and document these characteristics.
- The discussion touched on the challenges of capturing dynamic changes in models and the need for version control and effective dates.
- The importance of separating confidentiality and exposure of information in the AIBOMs was emphasized.
- The need to define the scope of the AIBOMs and focus on shippable models was suggested.
- The idea of placeholder categories for future use cases and dynamic models was proposed.
- The next steps include documenting the use cases in a template and prioritizing them for further work.
- Compliance was suggested as a high-priority use case to start with, followed by vulnerability and risk assessment.

Action items:

Everyone

- Write the use cases into the provided template
- Review the use cases in the next iteration
- Register for the AI workshop
- Start working on the compliance use case
- Prioritize the use cases in the following order: compliance, vulnerability, risk assessment, etc.

Meeting Date: Jul 22, 2024

Attendees

- Add yourself here
- Emily Fox (Red Hat)
- Allan Friedman (CISA)

- Daniel Bardenstein
- Helen Oakley
- Arthit Suriyawongkul
- John Nuckles
- Janane Suresh

Notes

- Agenda goal: get us to 5-8 initial use cases for Project 1
- In order to become an Official Tiger Team, we have to update/finalize our Tiger Team proposal/overview document with the following
 - Project name:
 - Lead(s) Name and Email:
 - Goal Output / Define "Done":
 - Why / Motivation:
 - In Scope:
 - Out of Scope:
 - Estimated Timeline / End Date:
 - Desired Skillsets / Expertise / Perspectives for Volunteers:
 - Project initial / scoping call time/link:
- Spent 10 minutes reading through use cases doc: https://docs.google.com/spreadsheets/d/1165T3oQvxOjKvcSbR-txmKAluXffRZ_75yZ_AjyeACxY/edit?gid=0#gid=0
- Live Q&A, discussions
 - Consolidate UC1,5,and 6 around model inventory and inventory-related use cases.
 - Incident Response?
 - Consolidate: UC4, 8, 10 around risk in third-party AI-enabled products
 - Expand IP use case into Fair Use
 - Consolidate 8 and 11 into Regulatory + Compliance use cases
- Criteria for final 5-8 use cases
 - Doesn't have to be 8 explicitly, just trying to narrow it down and consolidate
 - Clearly distinct for every other use case
 - Clear outcome of the business use case
- Questions/discussion
 -
 - How do we consider different roles & responsibilities in use cases?

Action Items

- Finalize Tiger Team doc with the above information (TODO)
- Add comments to UC doc (All)
- Consolidate UCs, present new list next week (TODO)

Meeting Date: Jul 15, 2024

Attendees

- Emily Fox (Red Hat)
- Daniel Bardenstein
- Amy Villasenor
- Allan Friedman
- Victoria Ontiveros
- Heimann Scott
- John Cavanaugh
- Bill Jacqmein
- Arthit Suriyawongkul
- Helen Oakley
- Marek Grac
- Ixchel Ruiz (Karakun)
- Ricardo Reyes
- Victor Lu
- Kim Sevenz
- John Nuckles
- Prosunjit Biswas
- Mark Orsi (GRF | BRC)
- Girish J (Splunk)
- Shafia Zubair
- Bunny Banowsky
- Ameya Naik
- Yotam Perkal
- Rhea Anthony
- Please

Notes

- Treasury Guidance, AI nutrition labels - Allan will follow back up with that team.
 - Under FSCC (financial sector coordinating council)
- [Use Case Brainstorming](#)
 - Built spreadsheet off the last session. Columns inform content needed
 - Note: Risks for model and risks for datasets may wish to be considered separate from each other, TBD
 - Alternatives are considerations for items that don't make sense to include in the AIBOM or may be on the fence for inclusion.
 - Look at leveraging the personas from SBOMs to shortcut the work:
 - defined four roles: Producer, Chooser, Operator, Subscriber.
 - Conducted collaborative review of the Use Cases

- UC-01, AI model inventory in a structured machine readable inventory.
 - We want an AI BOM so we can do vuln mgmt and lifecycle planning of the Models, and their components in use by my organization or that are included in the software I acquire.
 - - inventory for -> search (“where is the Stable-Diffusion@1.2.3” in my company)
 - - inventory for -> vuln mgmt (“where is \$vulnerableModel in my company.”)
 - - inventory for -> incident response (“what models were trained on LAION-5B in my environment “)
 - - inventory for -> data mgmt (“what data sources and their provenance contributed to my use”)
 - Recommendation: hold for now until we’ve reviewed others to collapse or expand as appropriate.
- UC-02, Assessing OS models risk
 - What real life examples does this group have on the persona responsible for assessing risk?
 - Reasonable amount of variability from examples shared as a result of internal organizational structures.
 - Persona: Any individual or team with the authority or responsibility for assessing the risk, surfacing risk, and/or providing a decision on that risk. It also considers anyone consuming that risk, internal or external.
 - Why: receive sufficient information to inform or necessitate additional compensating and mitigating controls or measures.
- The intent of the Use Case is to surface the information necessary to be present within the AIBOM.
- We may nominate or consider use cases to be consolidated or expanded.
- UC-08 - merging and sharing - higher level policy
- UC-11 - internal - organizational which may be more stringent
- Policy Updates
 - Post links and share!
 -
- AI Events & Conferences
 - Post links and share!
 - NeXt-generation Data Governance workshop 2024
 - co-located with 20th SEMANTiCS
 - 17 September 2024
 - Amsterdam, Netherlands
 - <https://nxdg-workshop.github.io/2024/>
 - From the CfP, there are topics in AI and data provenance
 - BH/B-Sides/DC

- AI Summit
- AI village
- Action Items
 - Tag yourself in the spreadsheet (Column C) for the content you want to own.
 - https://docs.google.com/spreadsheets/d/1165T3oQvxOjKvcSbR-txmKAluXffRZ_75yZAJyeACxY/edit?usp=sharing
 - As you fill in the cells please add your initials/name before the content you place in the cell, it will allow multiple individuals to reflect their perspective within the cell so we can refer back during discussion to arrive at consensus.
 - Example: (within 1 cell)
- EF: the objective for this use cases is....
- DB: the objective is broader and with these exceptions
- IH: The objective is to manage risks effectively in the absence of other information.

Action Items

- AI: ...

Meeting Date: Jul 1, 2024

Attendees

- Daniel Bardestein!
- Helen Oakley (SAP)
-

Notes

Started 2:05

.

- AI relevance for listed problems:

<https://github.com/aibom-squad/AIBOM-Tiger-Team/blob/main/AIBOM-problem-statements.md>

- Matt Rutkowski: If it is helpful, in a future call I can review the CDX model card data schema
- Nicholas Vidovich (Finite State): I think there is also an open question about how much overlap exists between model cards and the AI BOM, and whether they are represented independently or if AI BOM is a superset of model cards?

Meeting Date: Jul 1, 2024

Attendees

- (add your name here)

Agenda

- Survey
- AIBOM 101: slides [here](#)

Action Items

- Please fill out the Survey by COB Friday: <https://forms.gle/AT8DsHKCq9rMWgZf8>
-

Notes

- (text here)

Meeting Date: Jun 24, 2024

Attendees:

-
- ...

Intro: add comment in the chat

- Name, role, employer
- One thing you hope to learn / achieve with this group

Action Items:

- Projects for this group - consolidated list and voting/ priority
- Respond to the Google Form that is going to come out.
 - Soliciting additional input
 - Availability for participation
- Daniel will send out Slack invitations with existing channels to everyone.
- Helen will create a github repo to serve as a landing page for this group under:
 - <https://github.com/aibom-squad>

Notes:

- Agenda: [Slides](#)
- Avoid nomenclature at this time so we can move forward and get something accomplished!
- Scoping discussion
 - Application, viability, and operationalizing SBOMs for AI - software, models, and datasets
 - NOT everything under the AI sec or AI gov spaces. There are plenty of other works working on vulns, SBOMs, etc. that have their own coverage.
 - No dissent, adopt and move forward
- Defining the problems
 - What are the problems people are having in this space?
 - Education on AIBOMs and their importance
 - Conveyance information about models and datasets in AI-related software in existing SBOM
 - What does a (AI)BOM solve for?
 - What are those use cases?
 - What is our problem statement?
 - Systemic identification of data that could reduce risk/ensure confidence in models
 - Accuracy conveyance
 - Privacy
 - Exposure of sensitive data
 - Bias
 - Can an AIBOM provide sufficient information in potentially identifying bias within a model?
 - Provenance
 - Can the AIBOM provide sufficient provenance on the Data source?
 - Data
 - Classification of datasets, source and quality
 - Reasoning about the model
 - What information is needed to inform others about what is "in" the software using AI?
 - Resolving drift in models

- BOMS are point in time for version delivered
 - Governance regarding AIBOM use/reuse
 - Transparency
 - How the model arrived at output
- Properties and Elements of AIBOMs
 - Happening in SPDX and CycloneDX - if we identify additional elements, then we will facilitate this to those entities.
- Use Cases?
 - How do security teams manage risk of open source models and data sets?
 - Third party AI risk
 - AI/ML risk
 - Use of the model, or derivatives of those models
 - Incident response for AIML related issues
 - Compromise of dataset, models impacted, downstream use of those models in other applications, software, and decision making.
- Projects for this group (no more than two projects in parallel at a given time.)
 - AIBOM Overview Paper - what is it, how is it used, when is it produced, when and how it can be used.
 - Use Cases (support from the group on this need)
 - Discussion that this can inform and drive the Paper.
 - AIBOM Components and fields - most important fields that drive risk information or may be reasoned about for risk decision making
 - AIBOM representation examples
 - Automating AIBOM generation - how they are produced for accuracy (too much free text!, we need to identify opportunities to automate)
 - Summarization and Outcome Paper - post activities
- Future decisions
 - AI VEX - potentially for decision later=