

SPDX SBOMs: Enabling Automation of Safety & Security Analysis

Kate Stewart

VP Dependable Embedded Systems, The Linux Foundation

kstewart@linuxfoundation.org



Software is Used in Critical Systems Today



Chemical



Financial



Commercial
Facilities



Food &
Agriculture



Communications



Government
Facilities



Critical
Manufacturing




Healthcare &
Public Care



Dams



Information
Technology



Defense
Industrial
Base



Nuclear
Reactors,
Materials,
& Waste



Emergency
Services



Transportation
Systems



Energy



Water &
Wastewater
Systems

Critical Infrastructure Today: Mix of Open & Proprietary

98%

Percent of general codebases and Android apps that contained OSS

[Synopsys2021]

70%

Percent of codebase that was OSS on average

[Synopsys2020]

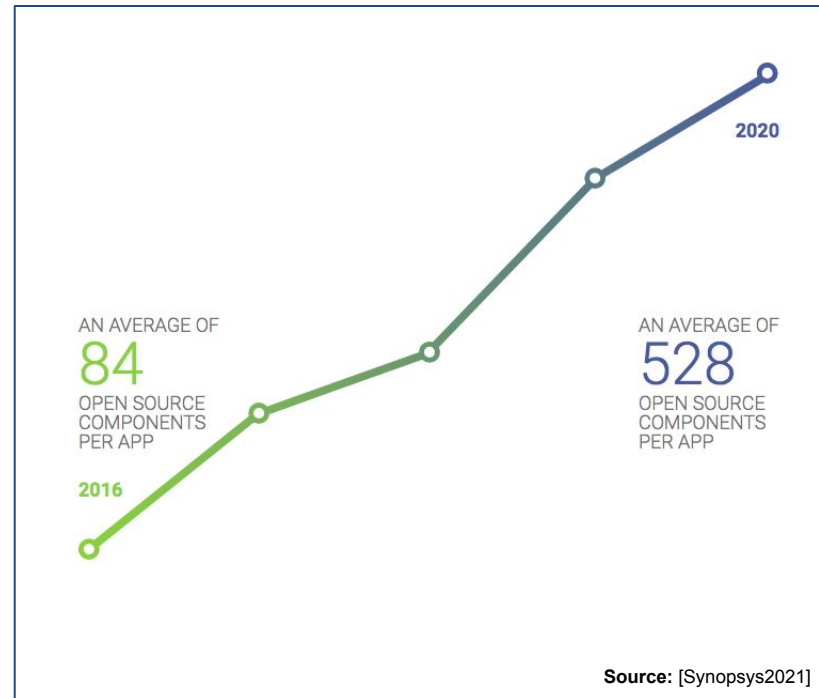
Source:

[Synopsys2020] "2020 Open Source Security and Risk Analysis Report" by Synopsys

<https://www.synopsys.com/content/dam/synopsys/sig-assets/reports/2020-ossra-report.pdf>

[Synopsys2021] "2021 Open Source Security and Risk Analysis Report" by Synopsys

<https://www.synopsys.com/software-integrity/resources/analyst-reports/open-source-security-risk-analysis.html>



Source: [Synopsys2021]

Cybersecurity & Critical Infrastructure

Critical Infrastructure

Since 2005, the 'Cybersecurity Policy for Critical Infrastructure Protection' has been set as a common action plan shared between the government, which bears responsibility for promoting independent measures by CI operators relating to CI cybersecurity and implementing other necessary measures, and CI operators which independently carry out relevant protective measures, and the new edition was published in 2022.

This document identifies the 14 sectors as critical infrastructure and it expects stakeholders to undertake the five measures as below.

1. Enhancement of Incident Response Capability
2. Maintenance and Promotion of the Safety Principles
3. Enhancement of Information Sharing System
4. Utilization of Risk Management
5. Enhancement of the Basis for CIP

2. Maintenance and promotion of the safety principles	Basically keep the element of "[1] Maintenance and promotion of the safety principles"	<ul style="list-style-type: none">◇ Clarify that safety standards, etc., that contribute to the enhancement of incident response capability and risk management are to be developed.◇ Consider survey methods capable of continuously improving the activities of CI operators.
---	--	--

The Cybersecurity Policy for Critical Infrastructure Protection

 Full Text

 [Guideline for Establishing Safety Principles for Ensuring Information Security of Critical Infrastructure\(5th Edition\)\(Revised on May 2019\)](#)

 [Risk Assessment Guide Based on the Concept of Mission Assurance in Critical Infrastructure \(1st Edition\)\(Revised on May 2019\)](#)

Maintenance and Promotion of Safety Principles

Safety Standards are looking for:

- **Unique ID**, something to uniquely identify the version of the software you are using.
 - Variations in releases make it important to be able to distinguish the exact version you are using.
 - The unique ID could be as simple as using the hash from a configuration management tool, so that you know whether it has changed.
- **Dependencies of the component**
 - Any chained dependencies that a component may require.
 - Any required and provided interfaces and shared resources used by the software component. A component can add demand for system-level resources that might not be accounted for.
- The component's **build configuration** (how it was built so that it can be duplicated in the future) and sources
- Any **existing bugs and their workarounds**
- **Documentation** for application manual for the component
 - The **intended use** of the software component
 - **Instructions** on how to **integrate** the software component correctly and **invoke it properly**
- **Requirements** for the software component
 - This should include the results of any testing to demonstrate requirements coverage
 - Coverage for nominal operating conditions and behavior in the case of failure
 - For highly safety critical requirements, test coverage should be in accordance with what the specification expects (e.g., Modified Condition/Decision Coverage (MC/DC) level code coverage)
 - Any safety requirements that might be violated if the included software performs incorrectly. This is specifically looking for failures in the included software that can cause the safety function to perform incorrectly. (This is referred to as a cascading failure.)
 - What the software might do under anomalous operating conditions (e.g., low memory or low available CPU)

Maintenance and Promotion of Safety Principles


Requirements are needed to know you're “**done**” after applying a patch:

- Need to be able to ensure you have compliance to the updated system requirements after applying a patch
- Given the rate of change and vulnerabilities, we need a way to make this automated, so it needs to be machine readable
- For each file patched, what requirements does it interact with, what tests need to be rerun to regenerate the evidence

Software Bill of Materials (SBOMs) today:

- Machine readable - Identities & Dependencies are part of the minimum definition
- SPDX SBOMs can also enables recording and connecting the sources, assessments, vulnerabilities & patches, build & calibration data, tests, requirements and evidence ⇒
path to automation

Common Understanding of “SBOM”



“An SBOM is a formal record containing the details and supply chain **relationships** of various **components** used in **building software**.”

These components, including libraries and modules, can be open source or proprietary, free or paid, and the data can be widely available or access-restricted.”

Source: NTIA’s [SBOM FAQ](#)



CYBERSECURITY
& INFRASTRUCTURE
SECURITY AGENCY



NTIA SBOM Guidance

Minimum Elements	
Data Fields	Document baseline information about each component that should be tracked: Supplier, Component Name, Version of the Component, Other Unique Identifiers, Dependency Relationship, Author of SBOM Data, and Timestamp.
Automation Support	Support automation, including via automatic generation and machine-readability to allow for scaling across the software ecosystem. Data formats used to generate and consume SBOMs include SPDX, CycloneDX, and SWID tags.
Practices and Processes	Define the operations of SBOM requests, generation and use including: Frequency, Depth, Known Unknowns, Distribution and Delivery, Access Control, and Accommodation of Mistakes.

Source: https://www.ntia.gov/files/ntia/publications/sbom_minimum_elements_report.pdf.

NTIA Software Bill Of Materials (SBOM) Guidance - Minimum Elements

Data Field	Description
Supplier Name	The name of an entity that creates, defines, and identifies components.
Component Name	Designation assigned to a unit of software defined by the original supplier.
Version of the Component	Identifier used by the supplier to specify a change in software from a previously identified version.
Other Unique Identifiers	Other identifiers that are used to identify a component, or serve as a look-up key for relevant databases.
Dependency Relationship	Characterizing the relationship that an upstream component X is included in software Y.
Author of SBOM Data	The name of the entity that creates the SBOM data for this component.
Timestamp	Record of the date and time of the SBOM data assembly.

SPDX 2.2 +
([ISO/IEC 5962:2021](https://www.iso.org/standard/72431.html))
supports all required minimum elements
(as well the optional that are mentioned in report)

Checker available at:
<https://github.com/spdx/ntia-conformance-checker>

When should an SBOM be created or consumed?

Safety and Security expect that Configuration Management (CM) information will be **maintained throughout the software lifecycle.**

SBOMs provide a mechanism to **track key artifacts and dependencies** as well as other useful CM information

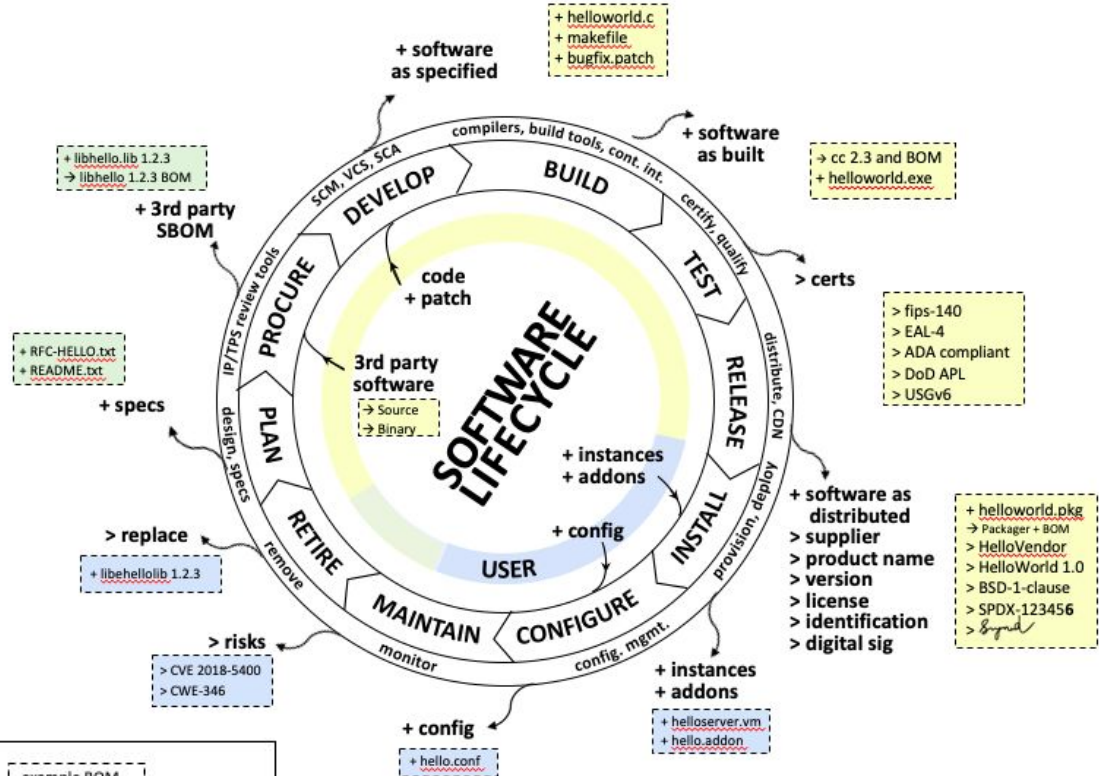


Image derived from : NTIA's Survey of Existing SBOM Formats and Standards

SBOM Types

SBOM TYPE	DEFINITION
Design	SBOM of intended, planned software project or product with included components (some of which may not yet exist) for a new software artifact.
Source	SBOM created directly from the development environment, source files, and included dependencies used to build an product artifact.
Build	SBOM generated as part of the process of building the software to create a releasable artifact (e.g., executable or package) from data such as source files, dependencies, built components, build process ephemeral data, and other SBOMs.
Deployed	SBOM provides an inventory of software that is present on a system. This may be an assembly of other SBOMs that combines analysis of configuration options, and examination of execution behavior in a (potentially simulated) deployment environment.
Runtime	BOM generated through instrumenting the system running the software, to capture only components present in the system, as well as external call-outs or dynamically loaded components. In some contexts, this may also be referred to as an “Instrumented” or “Dynamic” SBOM.
Analyzed	SBOM generated through analysis of artifacts (e.g., executables, packages, containers, and virtual machine images) after its build. Such analysis generally requires a variety of heuristics. In some contexts, this may also be referred to as a “3rd party” SBOM.

Source: [Types of Software Bills of Materials \(SBOM\)](#) published by CISA on 2023/4/21

KEY: Generate SBOMs **when** the data is available



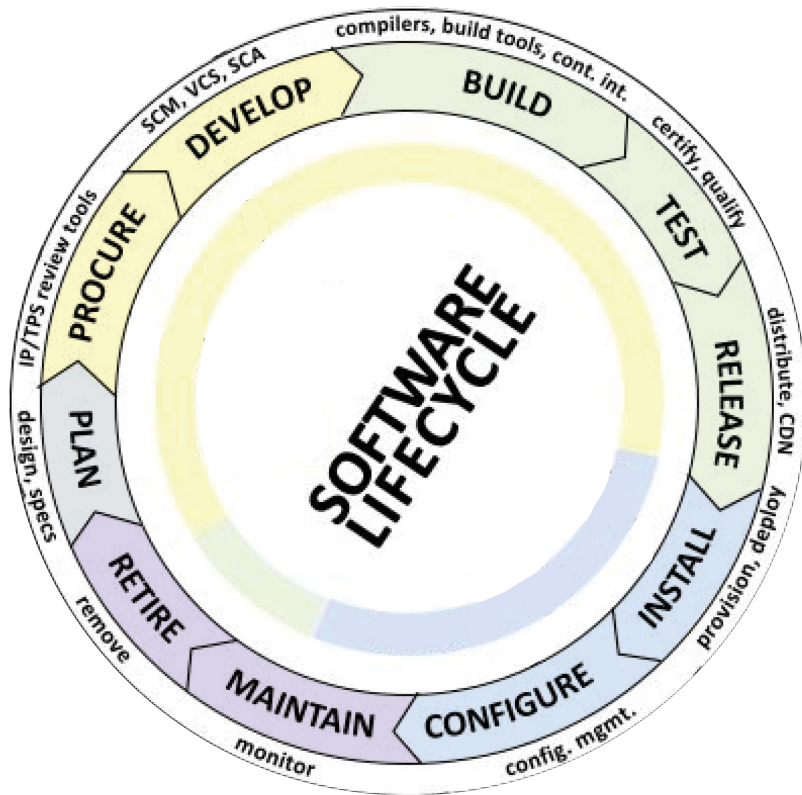
Source SBOM



Design SBOM



Runtime SBOM

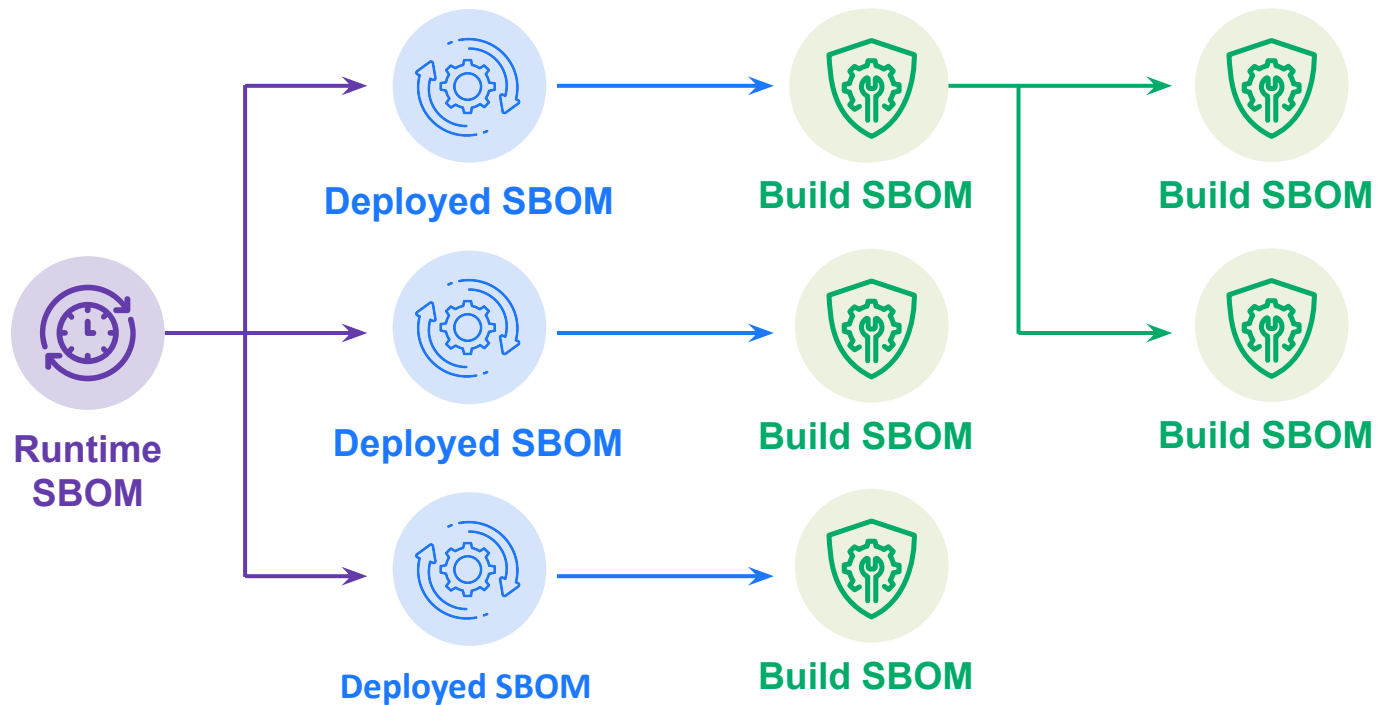


Build SBOM




Deployed SBOM

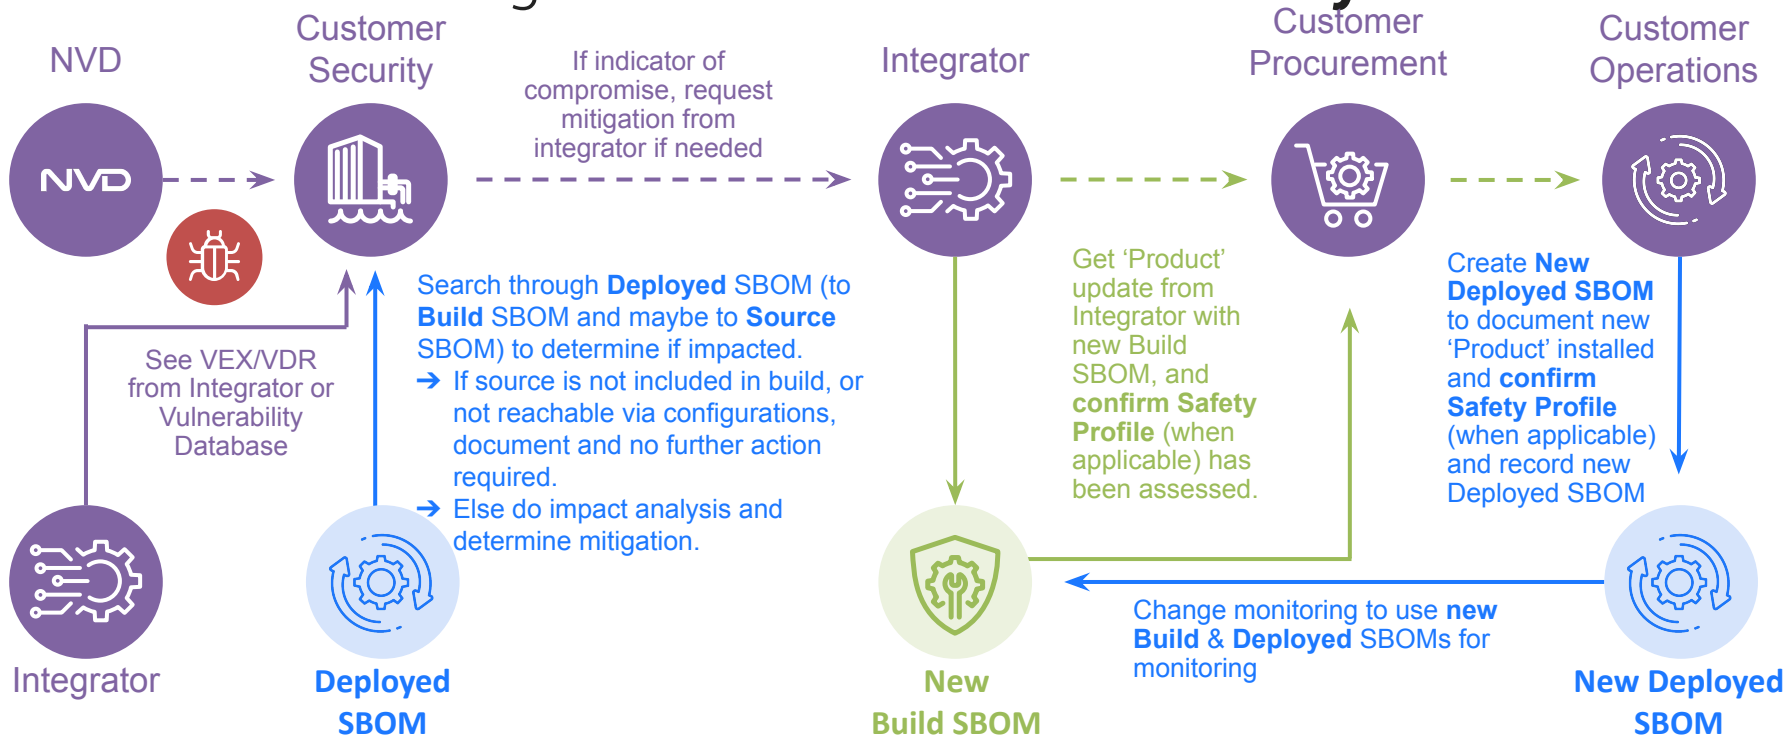
Understanding System: Traceability



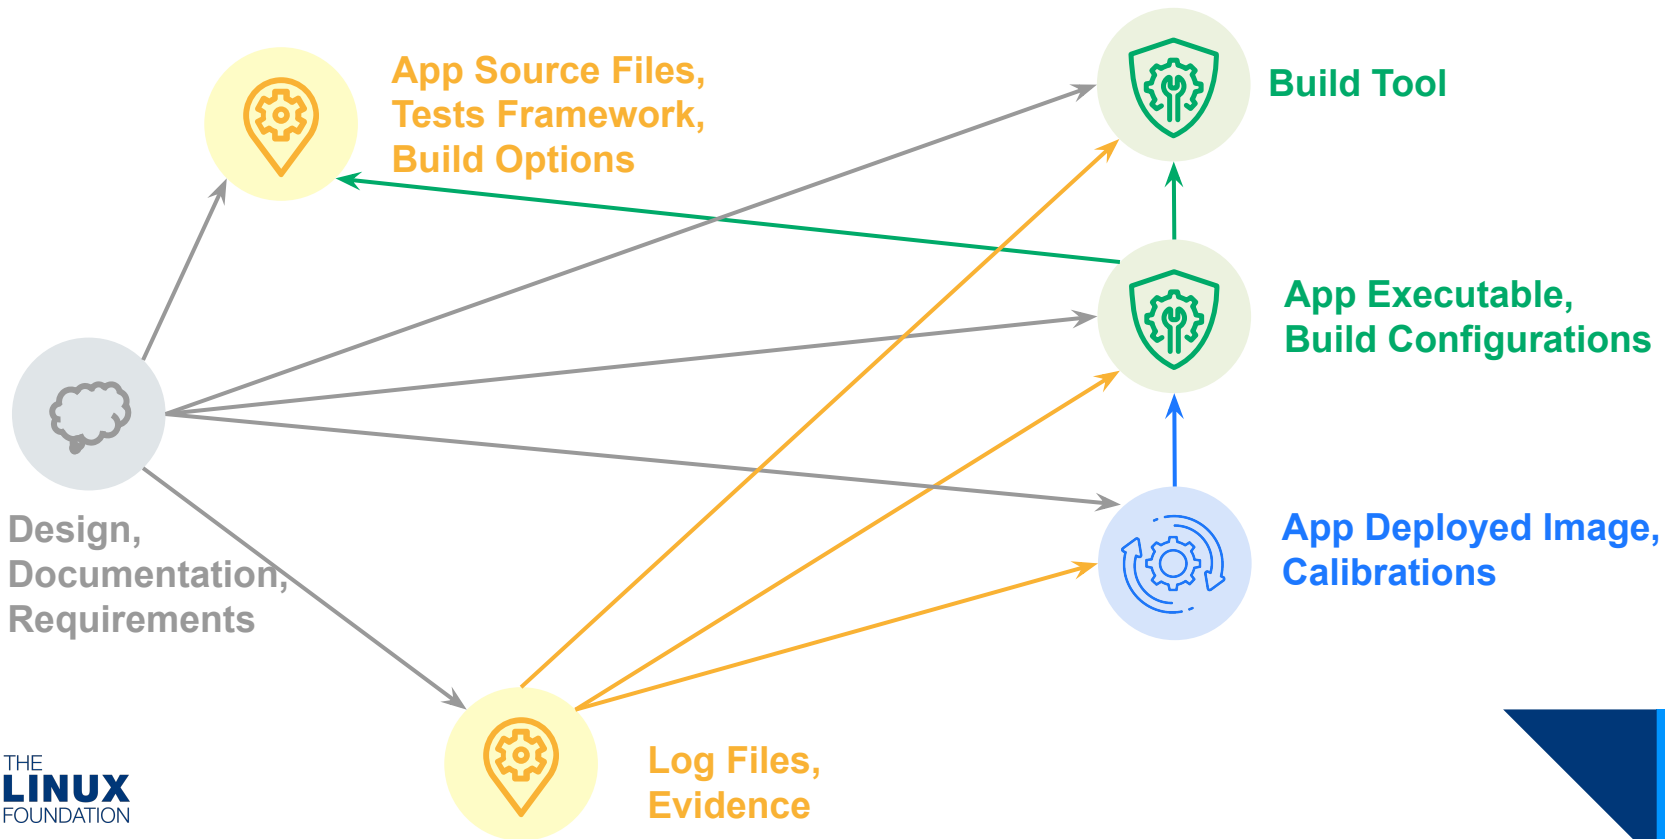
Understanding Safety Critical System: Traceability



Managing a security fix: Customer & Integrator need to **check Safety Profile**



SPDX SBOM's Enable Linking: Requirements to Code to Tests to Evidence



SPDX v2.3 Document must contain:

SPDX Document
Creation Information

SPDX v2.3 Document may contain:

Package Information

File Information

Snippet Information

Other Licensing
Information Detected

Relationships between SPDX
Elements Information

Annotations Information

Software Package Data Exchange (SPDX®) specification is a standard for communicating the component and metadata information associated with software

Charter: To create a set of **data exchange** standards that enable companies and organizations to share human-readable and machine-processable **software package metadata** to facilitate **software supply chain processes**.

ISO/IEC 5962:2021



- Able to represent SBOMs from binary images and track back to the source files and snippets.
- Specification is freely available from [ITTF site](#)
- Future updates are live tracked at: <https://spdx.github.io/spdx-spec> and work on satisfying safety requirements is being included
- More information at spdx.dev

A screenshot of the ISO website page for the ISO/IEC 5962:2021 standard. The page features a dark navigation bar at the top with a search icon, a shopping cart icon, the language "EN", and a "MENU" button. Below the navigation bar is the ISO logo and a red banner indicating the ICS number "ICS > 35 > 35.080". The main title is "ISO/IEC 5962:2021 Information technology – SPDX® Specification V2.2.1". A light blue box contains the text: "The electronic version of this International Standard can be downloaded from the ISO/IEC Information Technology Task Force (ITTF) web site." Below this is an "ABSTRACT" section with a "PREVIEW" button. The abstract text reads: "This Software Package Data Exchange® (SPDX®) specification defines a standard data format for communicating the component and metadata information associated with software packages. An SPDX document can be associated with a set of software packages, files or snippets and contains information about the software in the SPDX format described in this specification." At the bottom, there is a "GENERAL INFORMATION" section with a registered trademark symbol. The status is "© Published" and the publication date is "2021-08".

Source: <https://www.iso.org/standard/81870.html>
accessed on 2021/11/19



specification v2.3.0

v2.3

Search docs

Copyright

Introduction

Clause 1: Scope

Clause 2: Normative references

Clause 3: Terms and definitions

Clause 4: Conformance

Clause 5: Composition of an SPDX document

Clause 6: Document Creation Information

Clause 7: Package Information

Clause 8: File Information

Clause 9: Snippet Information

Clause 10: Other Licensing Information Detected

Clause 11: Relationship between SPDX Elements Information

Clause 12: Annotation Information

Clause 13: Review Information (deprecated)

Annex A: SPDX License List



»Copyright

The Software Package Data Exchange® (SPDX®) Specification Version 2.3

Copyright © 2010-2022 Linux Foundation and its Contributors. This work is licensed under the Creative Commons Attribution License 3.0 Unported (CC-BY-3.0) reproduced in its entirety in Annex J herein. All other rights are expressly reserved.

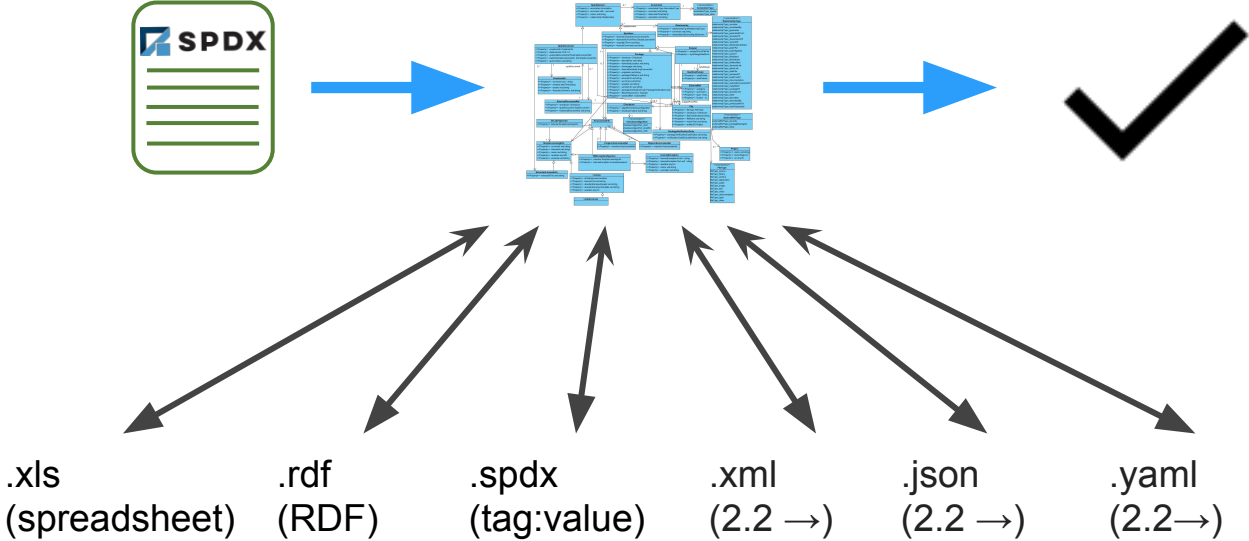
With thanks to Adam Cohn, Alan Tse, Alexios Zavras, Andrew Back, Ann Thornton, Basil Peace, Bill Schineller, Bradlee Edmondson, Bruno Cornec, Ciaran Farrell, Daniel German, David Edelsohn, David Kemp, David A. Wheeler, Debra McGlade, Dennis Clark, Dick Brooks, Ed Warnicke, Eran Strod, Eric Thomas, Esteban Rockett, Gary O'Neall, Guillaume Rousseau, Hassib Khanafer, Henk Birkholz, Hiroyuki Fukuchi, Jack Manbeck, Jaime Garcia, Jeff Luszcz, Jilayne Lovejoy, John Ellis, Jonas Oberg, Kamsang Salima, Karen Copenhaver, Kate Stewart, Kevin Mitchell, Kim Weins, Kirsten Newcomer, Kris Reeves, Liang Cao, Lon Hohberger, Marc-Etienne Vargenau, Mark Gisi, Marshall Clow, Martin Michlmayr, Martin von Willebrand, Mark Atwood, Matija Šuklje, Matt Germonprez, Maximilian Huber, Michael J. Herzog, Michel Ruffin, Nisha Kumar, Norio Kobota, Nuno Brito, Oliver Fendt, Paul Madick, Peter Williams, Phil Robb, Philip Koltun, Philip Odenice, Phillippe Ombredanne, Pierre Lapointe, Rana Rahal, Robert Martin, Robin Gandhi, Rose Judge, Sam Ellis, Sameer Ahmed, Scott K Peterson, Scott Lamons, Scott Sterling, Sean Barnum, Sebastian Crane, Shane Coughlan, Steve Cropper, Steve Winslow, Stuart Hughes, Thomas F. Incorvia, Thomas Steenberg, Tom Callaway, Tom Vidal, Toru Taima, Venkata Krishna, W. Trevor King, William Bartholomew, Yev Bronshteyn, Yoshiko Ouchi, Yoshiyuki Ito, Yuji Nomura and Zachary McFarland for their contributions and assistance.

source: <https://spdx.github.io/spdx-spec/v2.3/>

SPDX Continuously Improves

- **2010/02** - specification drafting began in a work-group of FOSSBazaar under Linux Foundation that came to be called "SPDX", was originally referred to as Package Facts.
- **2010/08** - "SPDX" announced as one of the pillars of the Linux Foundation's Open Compliance Program.
- **2011/08** - SPDX 1.0 specification - handles packages.
- **2012/08** - SPDX 1.1 specification - fixed flaw in verification algorithm
- **2013/10** - SPDX 1.2 specification - improved interaction with license list, additional fields for documenting project info.
- **2015/05** - SPDX 2.0 specification - added ability to handle multiple packages, relationships between packages and files, annotations.
- **2016/08** - SPDX 2.1 specification - added snippets, support for associating packages with external reference sources of information about packages, using SPDX License identifiers in files
- **2019/06** - SPDX 2.1.1 - conversion of specification from google docs to github as repository
- **2020/05** - SPDX 2.2 - Includes SPDX-lite
- **2020/08** - SPDX 2.2.1 prepared for submission to ISO.
- **2021/08** - ISO/IEC 5962 available
- **2022/08** - SPDX 2.3 published to improve interoperability with other formats
- **2023/Q2** - SPDX 3.0 release candidate and prototyping in progress ...

Formal Model Enables Validation & Interchange Between Specific File Formats



SPDX Relationships Clarify Dependency Types

DESCRIBES	DEPENDENCY_OF	PREREQUISITE_FOR	GENERATES	VARIANT_OF
DESCRIBED_BY	RUNTIME_DEPENDENCY_OF	HAS_PREREQUISITE	TEST_OF	FILE_ADDED
CONTAINS	BUILD_DEPENDENCY_OF	ANCESTOR_OF	TEST_TOOL_OF	FILE_DELETED
CONTAINED_BY	DEV_DEPENDENCY_OF	DESCENDENT_OF	TEST_CASE_OF	FILE_MODIFIED
DYNAMIC_LINK	OPTIONAL_DEPENDENCY_OF	DOCUMENTATION_OF	EXAMPLE_OF	PATCH_FOR
STATIC_LINK	PROVIDED_DEPENDENCY_OF	BUILD_TOOL_OF	METAFILE_OF	PATCH_APPLIED
AMENDS	TEST_DEPENDENCY_OF	EXPANDED_FROM_ARCHIVE	PACKAGE_OF	REQUIREMENT_FOR
COPY_OF	OPTIONAL_COMPONENT_OF	DISTRIBUTION_ARTIFACT	DATA_FILE_OF	SPECIFICATION_FOR
DEPENDS_ON	DEPENDENCY_MANIFEST_OF	GENERATED_FROM	DEV_TOOL_OF	OTHER

For more details see: <https://spdx.github.io/spdx-spec/v2.3/relationships-between-SPDX-elements/>



SPDX Generation Tooling

- [tools-java](#)
 - Aug 12, 2022 - [v1.1.0](#) update to support 2.3
 - Sept 2022 → Feb 2023 5 releases for performance improvements & fixes
- [tools-python](#)
 - OpenSSF Funded cleanup & restructuring
 - Dec 8, 2022 - [v0.7.0](#) - update to support 2.3 & clean up bug backlog
 - Prototyping of 3.0 in progress
- [tools-golang](#)
 - Jan 12, 2023 - [v0.4.0](#) - update to support 2.3
- [spx-online-tools](#) (validator & translator)
 - Aug 12, 2022 - [v1.0.7](#) add support SPDX 2.3
 - Nov 15, 2022 - [v1.0.9](#) add in [NTIA Conformance Checker](#) Tool

SPDX Consumption Tooling

- [spdx-online-tools](#) (validator & translator)
 - Aug 12, 2022 - [v1.0.7](#) add support SPDX 2.3
 - Nov 15, 2022 - [v1.0.9](#) add in [NTIA Conformance Checker](#) Tool
- [SPDX-to-OSV](#) (vulnerability lookup)
 - Produce an Open Source Vulnerability JSON file based on information in an SPDX document
 - Jan 10, 2022 - [v0.1.1](#) - pick up tooing updates
- [ntia-conformance-checker](#) (minimum SBOM fields present)
 - Check that an SBOM meets the minimum field requirements
 - Started as GSOC project, and maintainer from Chainguard has adopted
 - Feb 8, 2022 - v0.2.1 - fix NOASSERTION supplier case

Also worth looking at: <https://github.com/nyph-infosec/daggerboard>

SBOMs Everywhere in 2022...

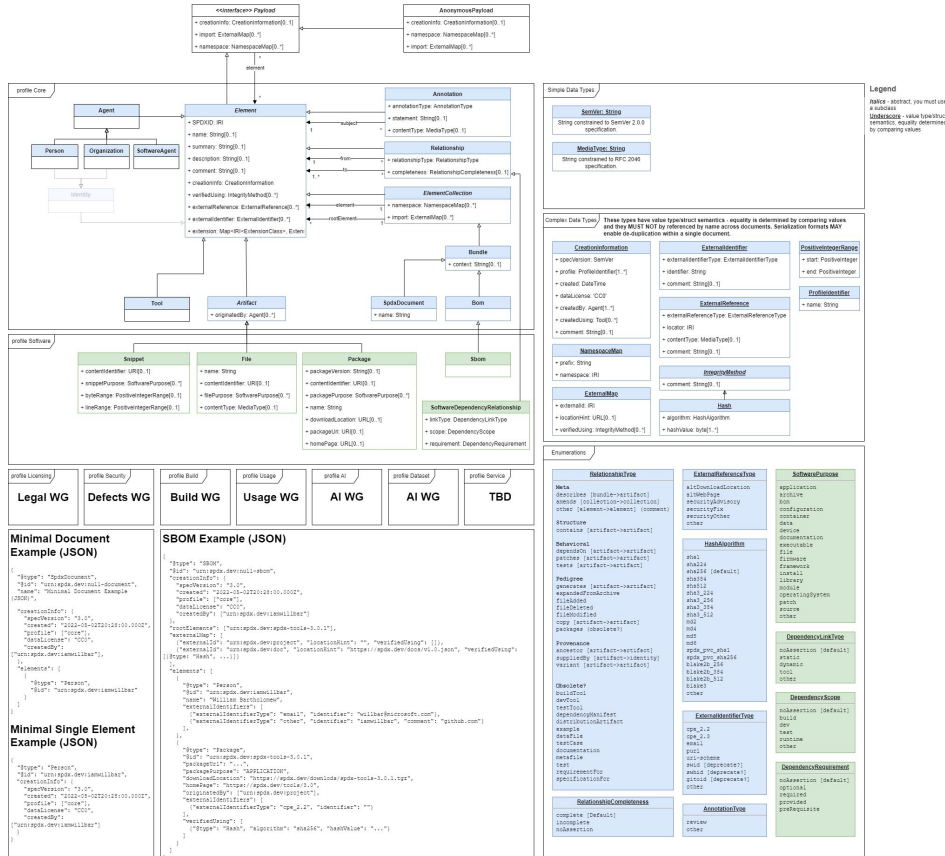
- OpenSSF - Work Stream 9 → **SBOM Everywhere SIG**
 - SPDX Python Library rework funded.
 - [Test suite](#) and [release candidate available now](#)
 - Started work on consolidating definitions of types of SBOMs → CISA working group to get broader adoption
 - Started documentation of use cases
- NTIA efforts have transferred to US DHS CISA
 - 4 working groups
 - SBOM Sharing, SBOM Adoption, SBOM Cloud, SBOM Tooling
 - International coordination with CERTs (like Japan's CERT) and other international government agencies

SPDX 3.0-rc1

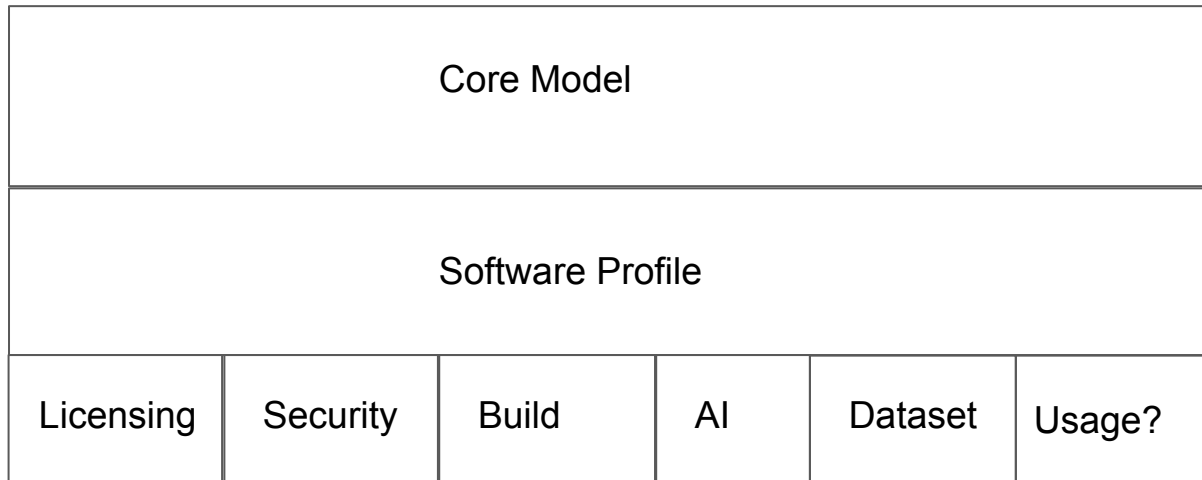
Why the Changes for SPDX 3.0?

- **Additional Use Cases**
 - AI and Data
 - Security and Defect information
- **Simplify**
 - Profiles
 - Remove confusing names
- **Flexibility**
 - Can communicate a single Element
 - Enhanced relationship structure with less relationship types (work in progress)

SPDX 3.0 Model

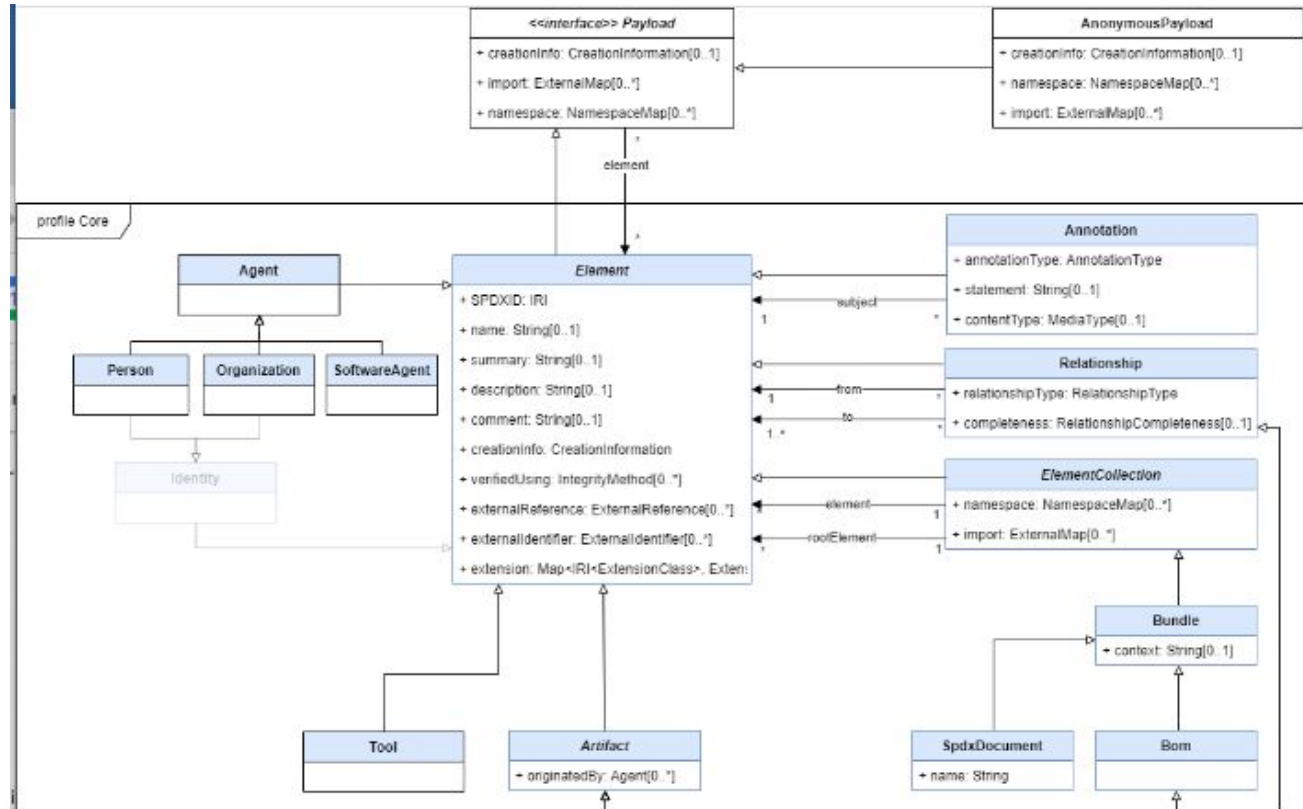


SPDX 3.0 - Increases Modularity

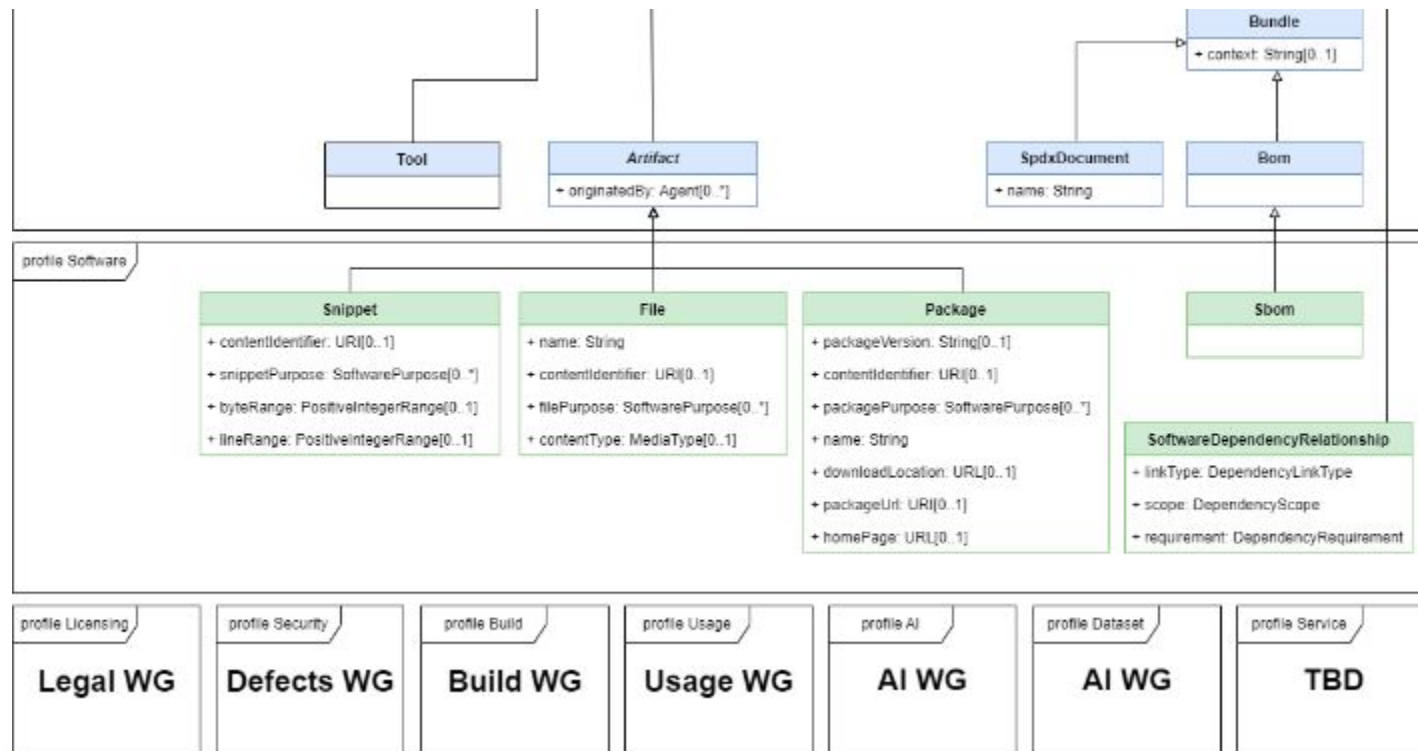


SPDX 3.0 (Core Model + Software Profile + Licensing Profile) == SPDX 2.3

SPDX 3.0 Core Model Permits Extensions



SPDX 3.0 Software Profile



SPDX 3.0 Specification Infrastructure

Specification is being transformed into markdown describing

- Classes, Properties, Enumerations
- Metadata (type & cardinality) and description for each element.
- Will be able to automatically generate schema from this version (for JSON, YAML, RDF, XML, tag-value, etc.) and reduce errors.

Profiles can add their own Classes and Properties and may also restrict other profiles (e.g. values, cardinalities, ...)

See: <https://github.com/spdx/spdx-3-model>

Licensing Profile Update

- Based on licensing-related fields in pre-existing SPDX spec, with updates:
 - More consistency across artifact types (package, file, snippet)
 - Aligning with SPDX 3.0 data model
 - Documenting the object model for license expressions
- Current status:
 - Fields were discussed in joint tech/legal calls in late 2020 and early 2021
 - Initial starting point draft shared in Mar. 2021
 - Revised earlier draft for new SPDX 3.0 model formatting
 - Initial draft included in main at:
<https://github.com/spdx/spdx-3-model/tree/main/model/Licensing>

For more information, contact: Steve Winslow or Alexios Zavras



Security Profile Update

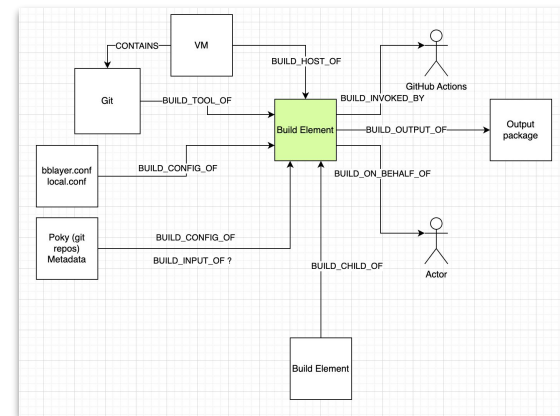
- Communicating vulnerabilities in software
 - Associate vulnerabilities with specific elements, like packages
 - Conveying vulnerability assessment (severity, impact, exploitability)
- Linking to external security information
 - *securityAdvisory*
 - Advisories and miscellaneous security related document
 - Common Security Advisory Framework (CSAF)
 - CycloneDX formatted security information
 - Open Source Vulnerability (OSV) document
 - Vulnerability Disclosure Report (VDR per NIST EO 14028)
 - *securityFix*
 - Code fix or patch for a security issue
 - *securityOther*
 - Any unspecified type of security information
- VEX support to assert status of a vulnerability
- Details at: <https://github.com/spdx/spdx-3-model/tree/main/model/Security>

For more information contact: Thomas Steenbergen, Jeff Schutt or Rose Judge

Build Profile Overview

Use Cases:

- Security
- Reproducibility
- Auditing quality/pedigree of build
- Safety
 - The source code at the time of release
 - The configuration used to build the software
 - The specific versions of the tools used to build the software



Producers: Build systems, secondarily, analysis tools

Initial draft has been submitted to model for discussion at :

<https://github.com/spdx/spdx-3-model/tree/main/model/Build>

Contact: Brandon Lumm or Nisha Kumar

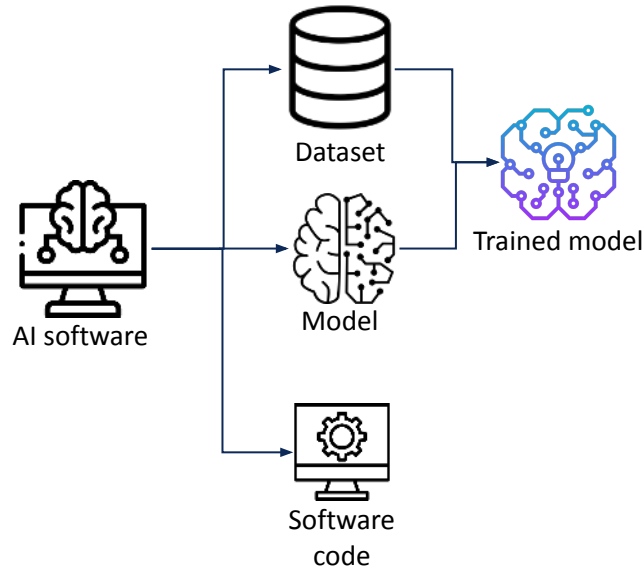
AI BOM \Rightarrow SPDX AI profile + SPDX Dataset profile

Traditional Software



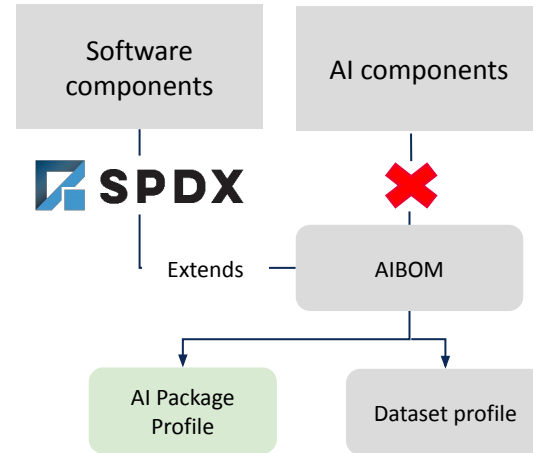
AI Software

Has additional elements that nuance that need to be captured to ensure its traceability



AI BOM

Enhances SPDX to describe AI software including and AI Software's components, licenses, copyrights, and security references.



Components of AI Application



Data

Schema

Sampling over Time

Volume

+



Model

Algorithms

More Training

Experiments

+



Code

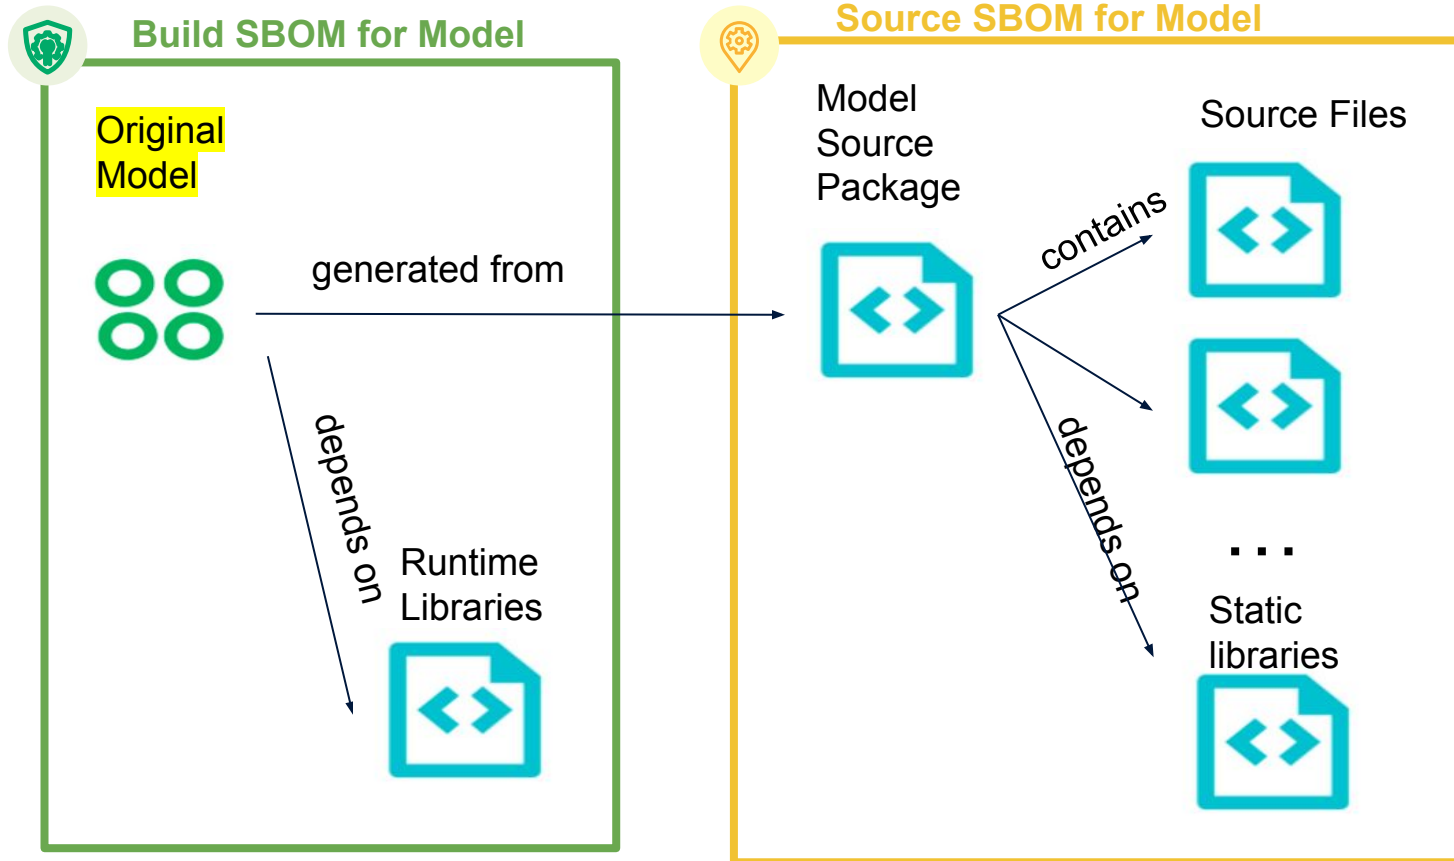
Business Needs

Bug Fixes

Configuration

Source: IBM Data Science Best Practices
(<https://ibm.github.io/data-science-best-practices/versioning.html>)

Build Up Original Model



Building the AI Application

Build SBOM for AI Application

AI Application Executable



Generated from


depends on

Trained Model



Source SBOM for AI Application

AI Application Source Package



contains

Source Files

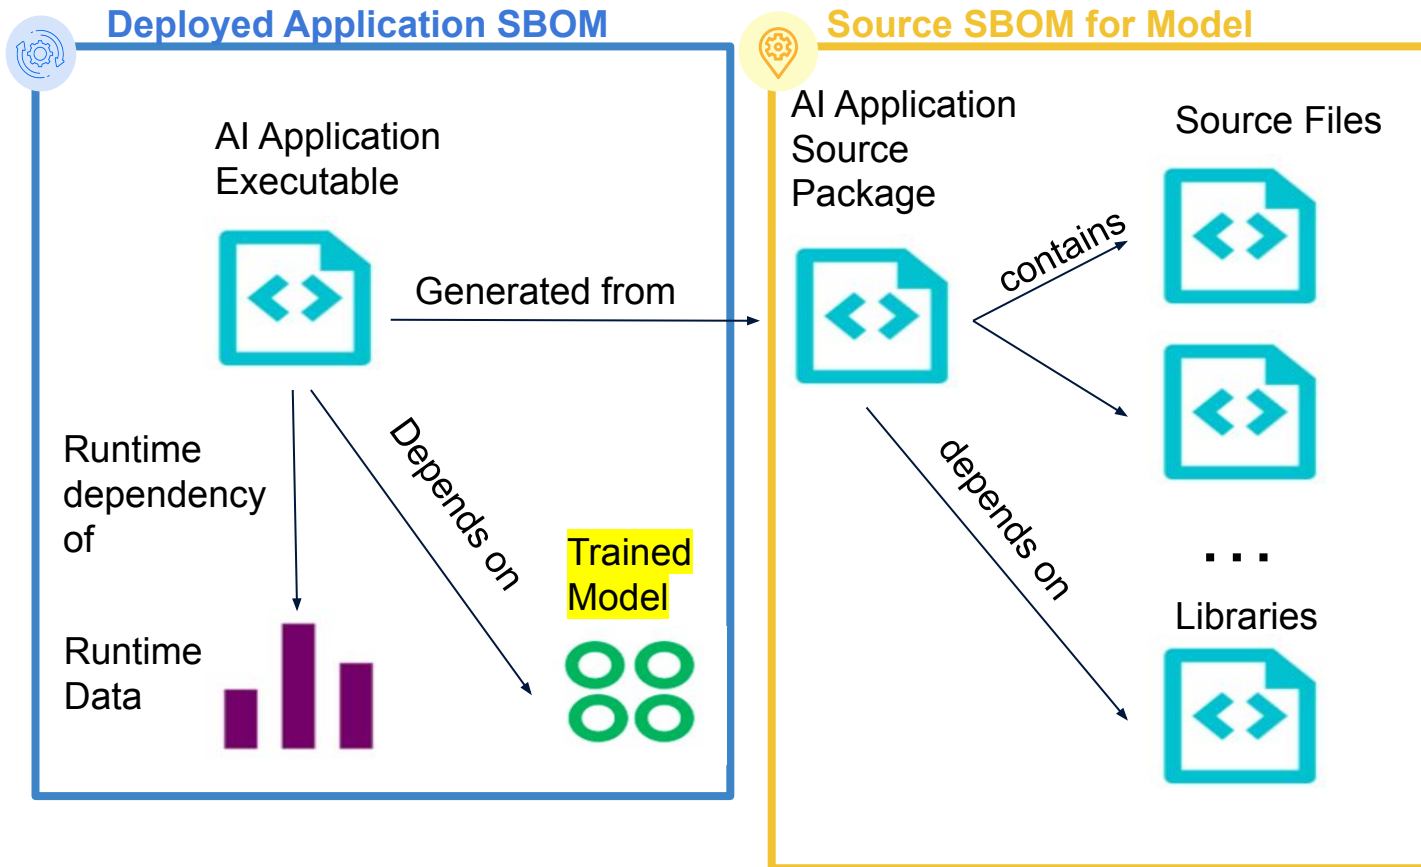


...

Libraries



Deploying the AI Application



SPDX can represent most of this today.

But what about the trained model?

GAP: Representing Training the model

Build SBOM for Trained Model



Trained
Model



Original
Model



Validation Data



Training Data



?

?

?

AI BOM Transparency Survey: Fields Required?

Datasheets

Datasheets for Datasets

Timnit Gebru¹ Jamie Morgenstern² Briana Vecchione¹ Jennifer Wortman Vaughan¹ Hanna Wallach¹ Hal Daume III^{1,4} Kate Crawford^{1,5}

Abstract


The machine learning community has no standardized way to document how and why a dataset was created, what information it contains, what tasks it should and should not be used for, and whether it might raise any ethical or legal concerns. To address this gap, we propose the concept of datasheets for datasets. In the electronics industry, it is standard to accompany every component with a datasheet providing standard operating characteristics, test results, recommended usage, and other information. Similarly, we recommend that every dataset be accompanied with a datasheet documenting its creation, composition, intended uses, maintenance, and other properties. Datasheets for datasets will facilitate better communication between dataset creators and users, and encourage the machine learning community to prioritize transparency and accountability.

1. Introduction

Machine learning is no longer a purely academic discipline. Domain such as criminal justice (Gervie et al., 2016; Systems, 2017; Andrews et al., 2006), hiring and employment (Mass & O’Neil, 2016), critical infrastructure (O’Connor, 2017; Choi, 2017), and finance (Lu, 2012), all increasingly depend on machine learning methods. By definition, machine learning models are trained using data, the choice of data fundamentally influences a model’s behavior. However, there is no standardized way to document how and why a dataset was created, what information it contains, what tasks it should and shouldn’t be used for, and whether it might raise any ethical or legal concerns. This lack of documentation is especially problematic when datasets are used to train models for high-stakes applications.

¹Microsoft Research, New York, NY ²Georgia Institute of Technology, Atlanta, GA ³Cornell University, Ithaca, NY ⁴University of Maryland, College Park, MD ⁵AI Now Institute, New York, NY. Correspondence to: Timnit.Gebru@microsoft.com.

Proceedings of the 37th Workshop on Fairness, Accountability, and Transparency in Machine Learning, Stockholm, Sweden, PMLR 80, 2018. Copyright 2018 by the author(s).



Model Cards

Model Cards for Model Reporting

Margaret Mitchell, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, Timnit Gebru
jmmitchelai,simonewu,andrewzaldivar,parkerbarnes,lucyvasserman,benhutchinson,elenaspitzer,timitgebru@protonmail.com

ABSTRACT

Trained machine learning models are increasingly used to perform high-impact tasks in areas such as law enforcement, medicine, education, and employment. In order to clarify the intended use cases of machine learning models and minimize their usage in contexts for which they are not well suited, we recommend that released models be accompanied by documentation detailing their performance characteristics. In this paper, we propose a framework that we call model cards, to encourage such transparent model reporting. Model cards are short documents accompanying trained machine learning models that provide benchmarked evaluation in a variety of conditions, such as across different cultural, demographic, or phenotype groups (e.g. race, geographic location, sex, Fitzpatrick skin type [15]) and intersectional groups (e.g. age and race, or sex and Fitzpatrick skin type) that are relevant to the intended application domains. Model cards also describe the context in which models are intended to be used, details of the performance evaluation procedure, and other relevant information. While we focus primarily on human-centered machine learning models in the application fields of computer vision and natural language processing, this framework can be used to document any trained machine learning model. To solidify the concept, we provide cards for two supervised models: One trained to detect smiling faces in images, and one trained to detect text-class content in text. We propose model cards as a step towards the responsible documentation of machine learning and related artificial intelligence technology increasing transparency into how well artificial intelligence technology works. We hope this work encourages those releasing trained machine learning models to accompany model releases with similar detailed evaluation numbers and other relevant documentation.

CCS NUMBERS

• General and reference — Evaluation; • Social and professional topics — User characteristics; • Software and its engineering — Use cases; Documentation; Software and its engineering — Human-centered computing — Weightthrough evaluation.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted by ACM for non-profit educational institutions registered with the ACM on the first page. Copyright for components of this work owned by others than ACM is held by their respective owners. For all other use, permission should be sought from the ACM or the copyright owner. This work is derived from arXiv:1808.07261v2 [cs.LG], January 20, 2018, Atlanta, GA, USA. ACM ISBN 978-1-4503-6250-5/18/01...\$15.00. http://dx.doi.org/10.1145/3187200.3187205.

arXiv:1810.03993v2 [cs.LG] 14 Jan 2019

arXiv:1803.09010v3 [cs.DB] 9 Jul 2018

FactSheets

FactSheets: Increasing Trust in AI Services through Supplier’s Declarations of Conformity

M. Arnold,¹ R. K. E. Bellamy,¹ M. Hind,¹ S. Houde,¹ S. Mehta,² A. Mojsilović,¹ R. Nair,¹ K. Natesan Ramamurthy,¹ D. Reimer,¹ A. Ohtamu,¹ D. Piorowski,¹ J. Tsuy,¹ and K. R. Varshney,¹ IBM Research

¹Yorktown Heights, New York, ²Bangaluru, Karnataka

Abstract

Accuracy is an important concern for suppliers of artificial intelligence (AI) services, but considerations beyond accuracy, such as safety (which includes fairness and explainability), security, and provenance, are also critical elements to engender consumers’ trust in a service. Many industries use transparent, standardized, but often not legally required documents called supplier’s declarations of conformity (SDCo) to describe the lineage of a product along with the safety and performance testing it has undergone. SDCo’s may be considered multi-dimensional fact sheets that capture and quantify various aspects of the product and its development to make it worthy of consumers’ trust. Inspired by this practice, we propose FactSheets to help increase trust in AI services. We envision such documents to contain provenance, performance, safety, security, and provenance information to be completed by AI service providers for examination by consumers. We suggest a comprehensive set of declarer items tailored to AI and provide examples for two fictitious AI services in the appendix of the paper.

1 Introduction

Artificial intelligence (AI) services, such as those containing predictive models trained through machine learning, are increasingly key pieces of products and decision-making workflows. A service is a function or application accessed two a customer via a cloud computing interface, typically by means of an application programming interface (API). For example, an AI ser-

vice could take an audio waveform as input and return a transcript of what was spoken as output, with all complexity hidden from the user, all computation done in the cloud, and all models used to produce the output pre-trained by the supplier of the service. A second more complex example would provide an audio waveform translated into a different language as output. The second example illustrates that a service can be made up of many different models (speech recognition, language translation, possibly sentiment or tone analysis, and speech synthesis) and is thus a distinct concept from a single pre-trained machine learning model or library.

In many different application domains today, AI services are achieving impressive accuracy. In certain areas, high accuracy alone may be sufficient, but deployments of AI in high-stakes decisions, such as credit applications, judicial decisions, and medical recommendations, require greater trust in AI services. Although there is no scholarly consensus on the specific traits that induce trustworthiness in people or algorithms [1, 2], fairness, explainability, general safety, security, and transparency are some of the issues that have raised public concern about trusting AI and threatened the further adoption of AI beyond low-stakes uses [3, 4]. Despite active research and development, however, there is no mechanism in place for the creator of an AI service to communicate how they are addressed in a deployed version. This is a major impediment to broad AI adoption.

Toward transparency for developing trust, we propose a FactSheet for AI Services. A FactSheet will contain sections on all relevant attributes of an AI service, such as intended use, performance, safety, security, and provenance. Performance will include appropriate accuracy or risk measures along with timing information. Safety, discussed in [5, 3] as the minimiza-

¹A. Ohtamu’s work was done while at IBM Research. Author is currently affiliated with Microsoft Research.

Source: <https://arxiv.org/pdf/1810.03993.pdf>

Source: <https://www.microsoft.com/en-us/research/uploads/prod/2019/01/1803.09010.pdf>

Source: <https://arxiv.org/pdf/1808.07261.pdf>

AI Profile Properties

AI/ML
Application



Model



Data



Required Fields:

- **Creator**
- **Supplier**
- **PackageVersionInfo**
- **DownloadLocation**
- **PackageDescription**
- **LicenseConcluded**
- **LicenseDeclared**
- **ReleaseTime**

Optional Fields:

- **Originator**
- **Checksum**
- **ValidUntilTime**
- **BuildTime**
- **PackageComments**
- *SensitivePersonalInformation*
- *EnergyConsumption*
- *StandardsCompliance*
- *InformationAboutTraining*
- *Hyperparameters*
- *SafetyRiskAssessment*
- *DataPreprocessingSteps*
- *ModelExplainabilityMechanisms*
- *MetricsDecisionThresholds*
- *Metrics*
- *Autonomy*
- *Domain*
- *Limitations*
- *Type*

Dataset Profile Properties

AI/ML
Application



Model



Data



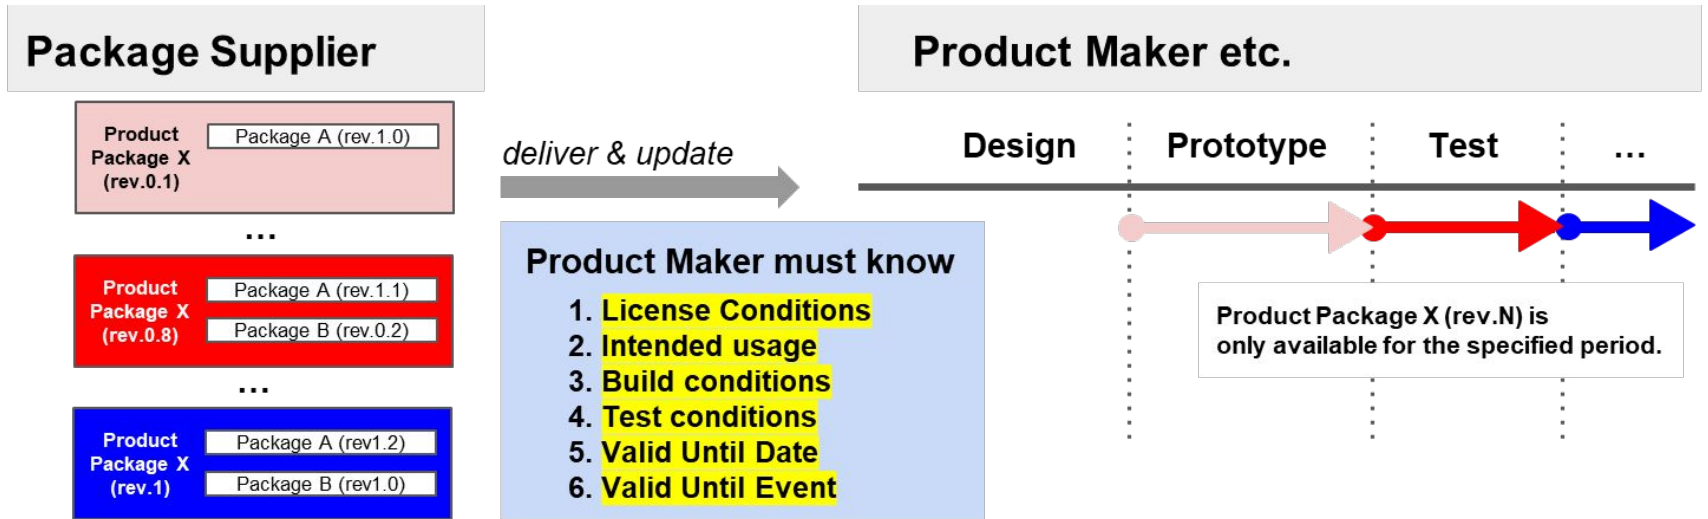
Required Fields:

- **Name**
- **Originator**
- **DownloadLocation**
- **LicenseConcluded**
- **LicenseDeclared**
- **PackageDescription**
- **BuiltTime**
- **ReleaseTime**
- *DatasetType*

Optional Fields:

- **Supplier**
- **VersionInfo**
- **Checksum**
- **ValidUntilTime**
- *IntendedUse*
- *DatasetCollectionProcess*
- *DatasetUpdateMechanism*
- *DatasetSize*
- *DatasetNoise*
- *KnownBias*
- *Errata*
- *SensorsUsed*
- *StandardCompliance*
- *SensitivePersonalInformation*
- *ConfidentialityLevel*
- *AnonymizationMethodUsed*

Usage Profile: to tell intentions as “Usage” for Delivery Product



Both parties must control the **Terms and Conditions** of using of the Delivery



What's next after After 3.0?

Future Direction: Hardware

- Safety Standards expect to know “system” that software is running on
- Vulnerabilities come from interaction between hardware and software (ie.heartbleed)
- Potential participants:
 - RISC-V & ARM core adopters,
 - Chips Alliance Members
 - Board Manufacturers
- For more information, contact: Kate Stewart

Future Direction: Safety Standards Automation

- Safety Standards expect to know
 - The source code at the time of production release
 - The documentation associated with the code
 - The configuration used to build the production software
 - The specific versions of the tools used to build the software
- Safety Standards Configuration Management (CM) Requirements are greatly simplified by following an effective SBOM process.
 - An SBOM supports **capturing the details** of what is in a specific release and supports determining what went wrong if a failure occurs.
 - The goal is to be able to **rebuild exactly** what the executable or binary was at the time of release.
- To learn more, see:
 - <https://www.linux.com/featured/sboms-supporting-safety-critical-software/>



Leverage SPDX Relationships to Support Safety Analysis



Using SPDX Relationship Information

Assumption: process to create and maintain all artefacts (requirements, architecture, tests, analysis report) is accepted and applied

Still the biggest pain: Keeping a complete and consistent set of documentation and verifying that the evidences are complete and consistent

SPDX style solution: Create SPDX Relationships between all documentation artefacts to track all possible system combinations!



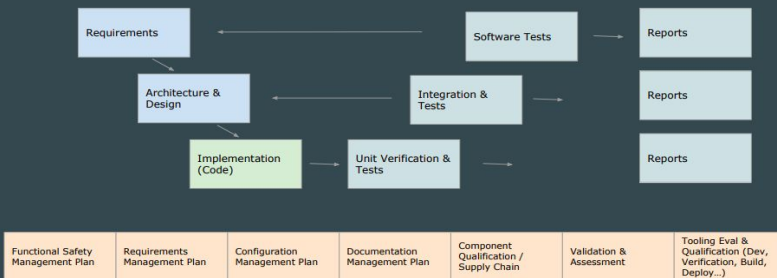
SPDX for product documentation

What kind of product documentation do we need to manage?

Specifications, Reports, Tests...

Safety Requirement Specification	a SPECIFICATION for functional requirements, architectural elements etc.
Unit Test	the TEST_CASE related to code or a specification artefact
Unit Test Report	DOCUMENTATION of a unit test EVIDENCE all tests have been performed as planned
Code	usually is GENERATED from or according to some specification artefact
Coding Guidelines	SPECIFICATION about the project specific details for the code

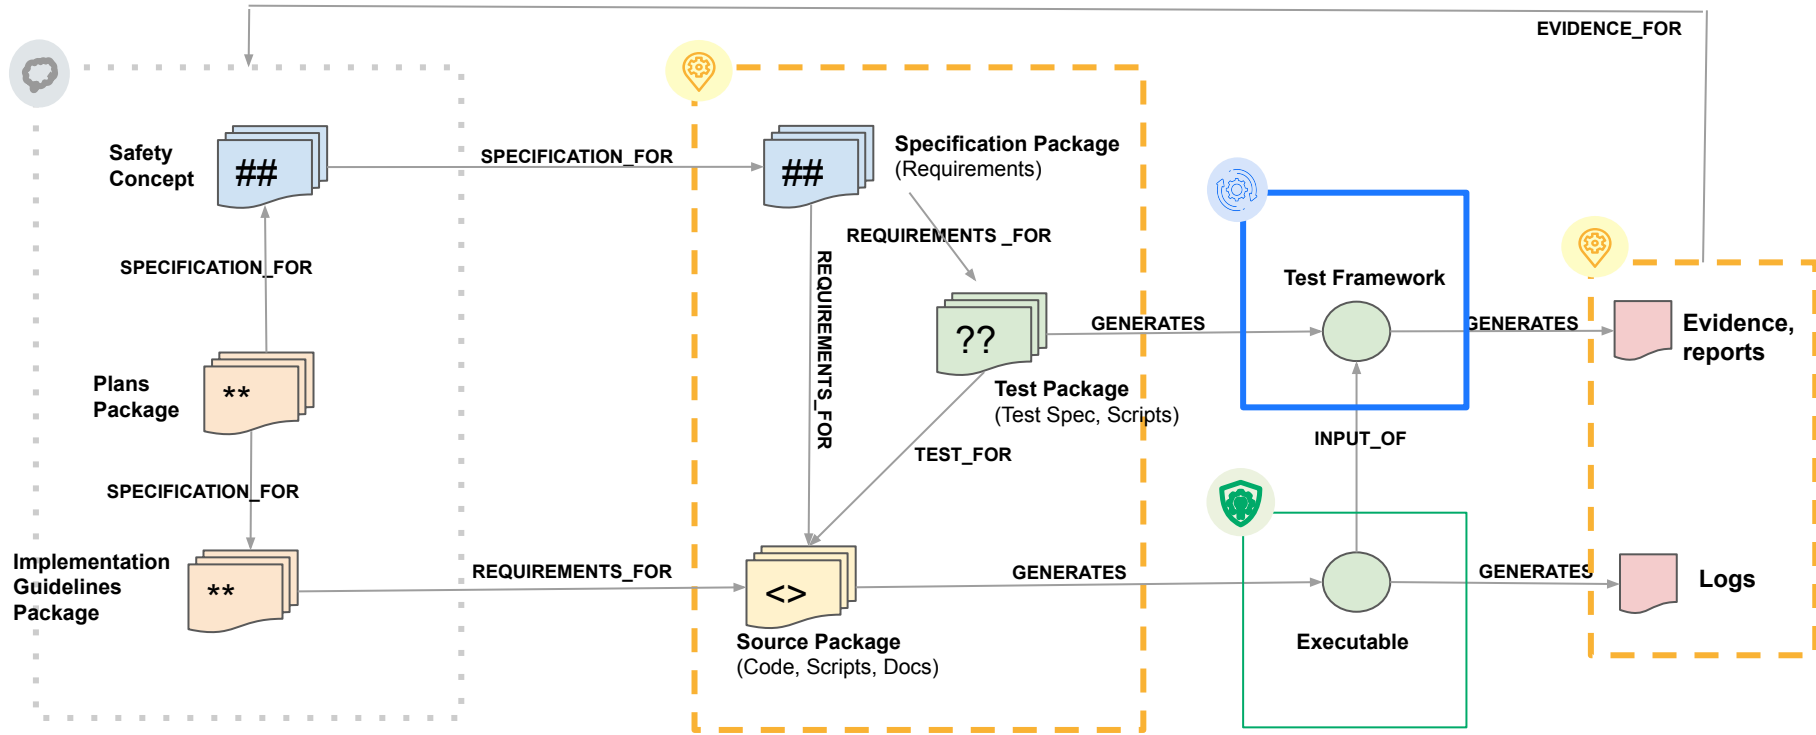
V-Model style documentation model



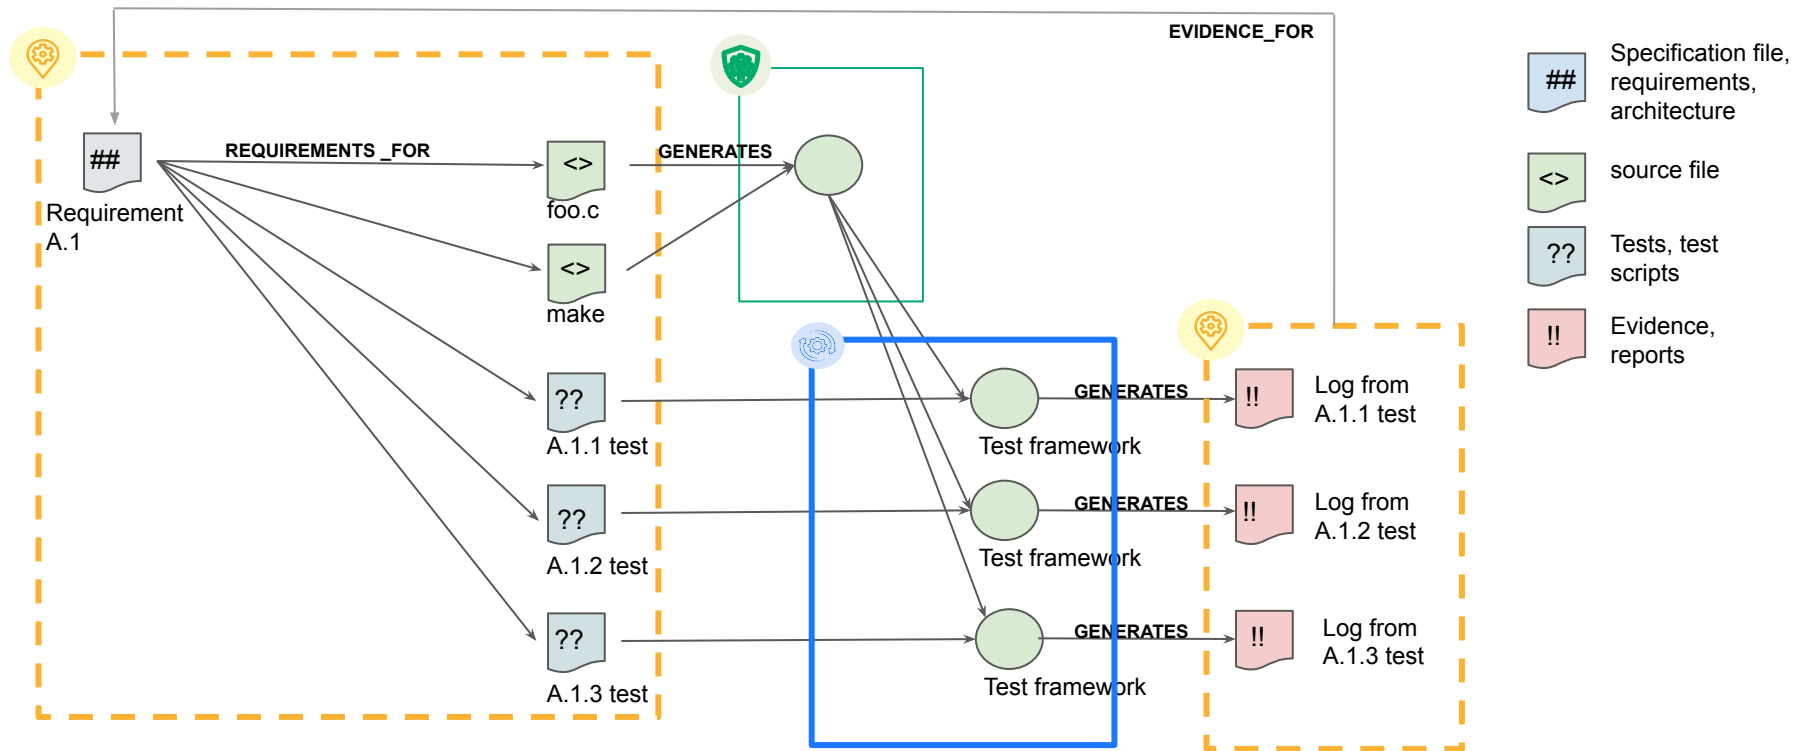
Source: https://fosdem.org/2023/schedule/event/sbom_fusa/



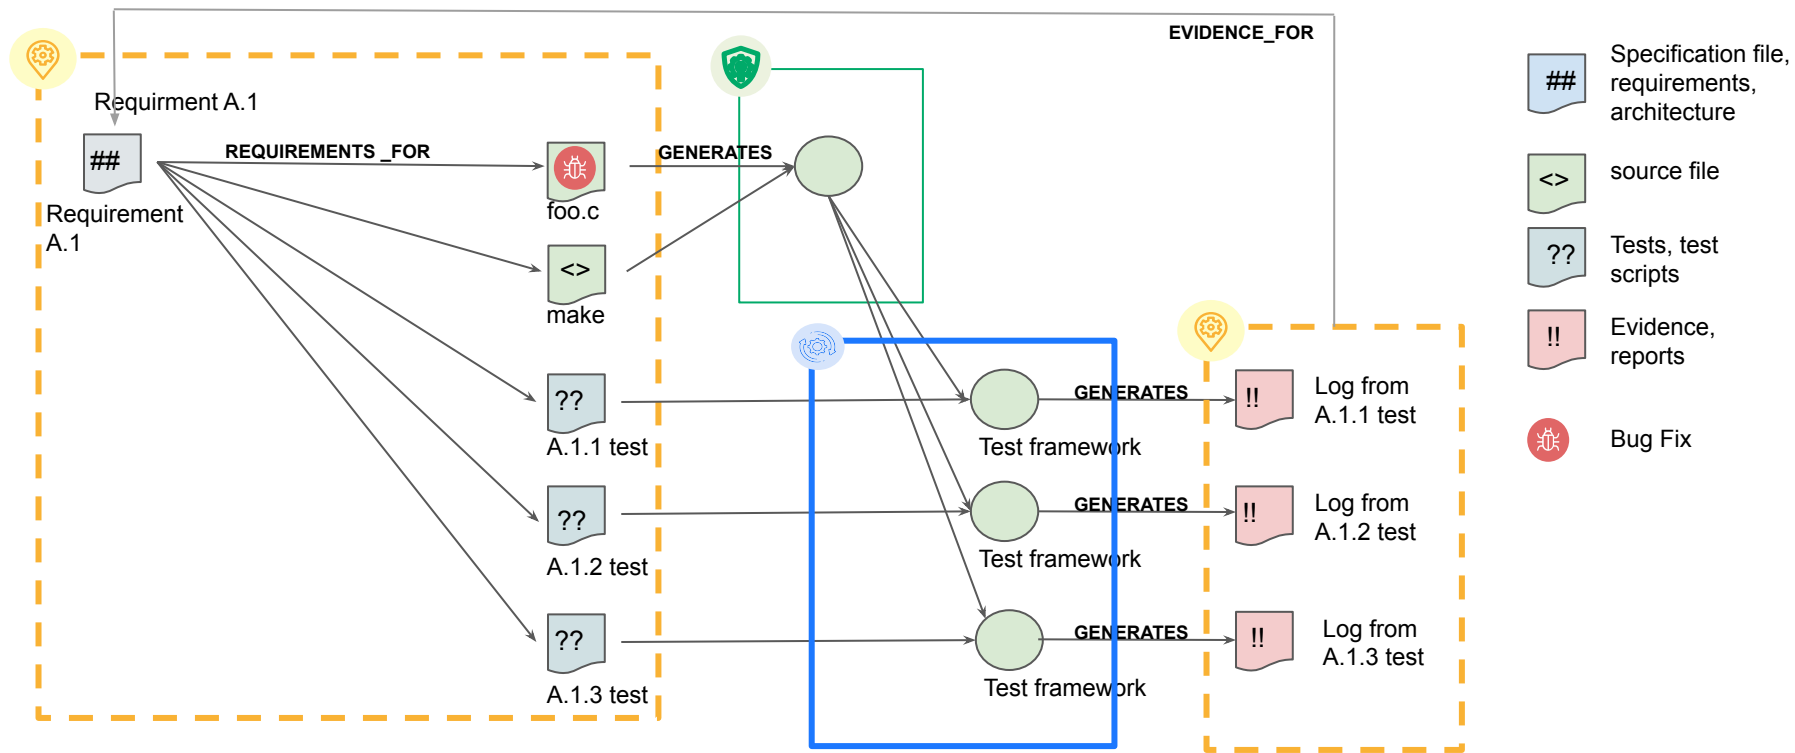
Leveraging SPDX Relationships



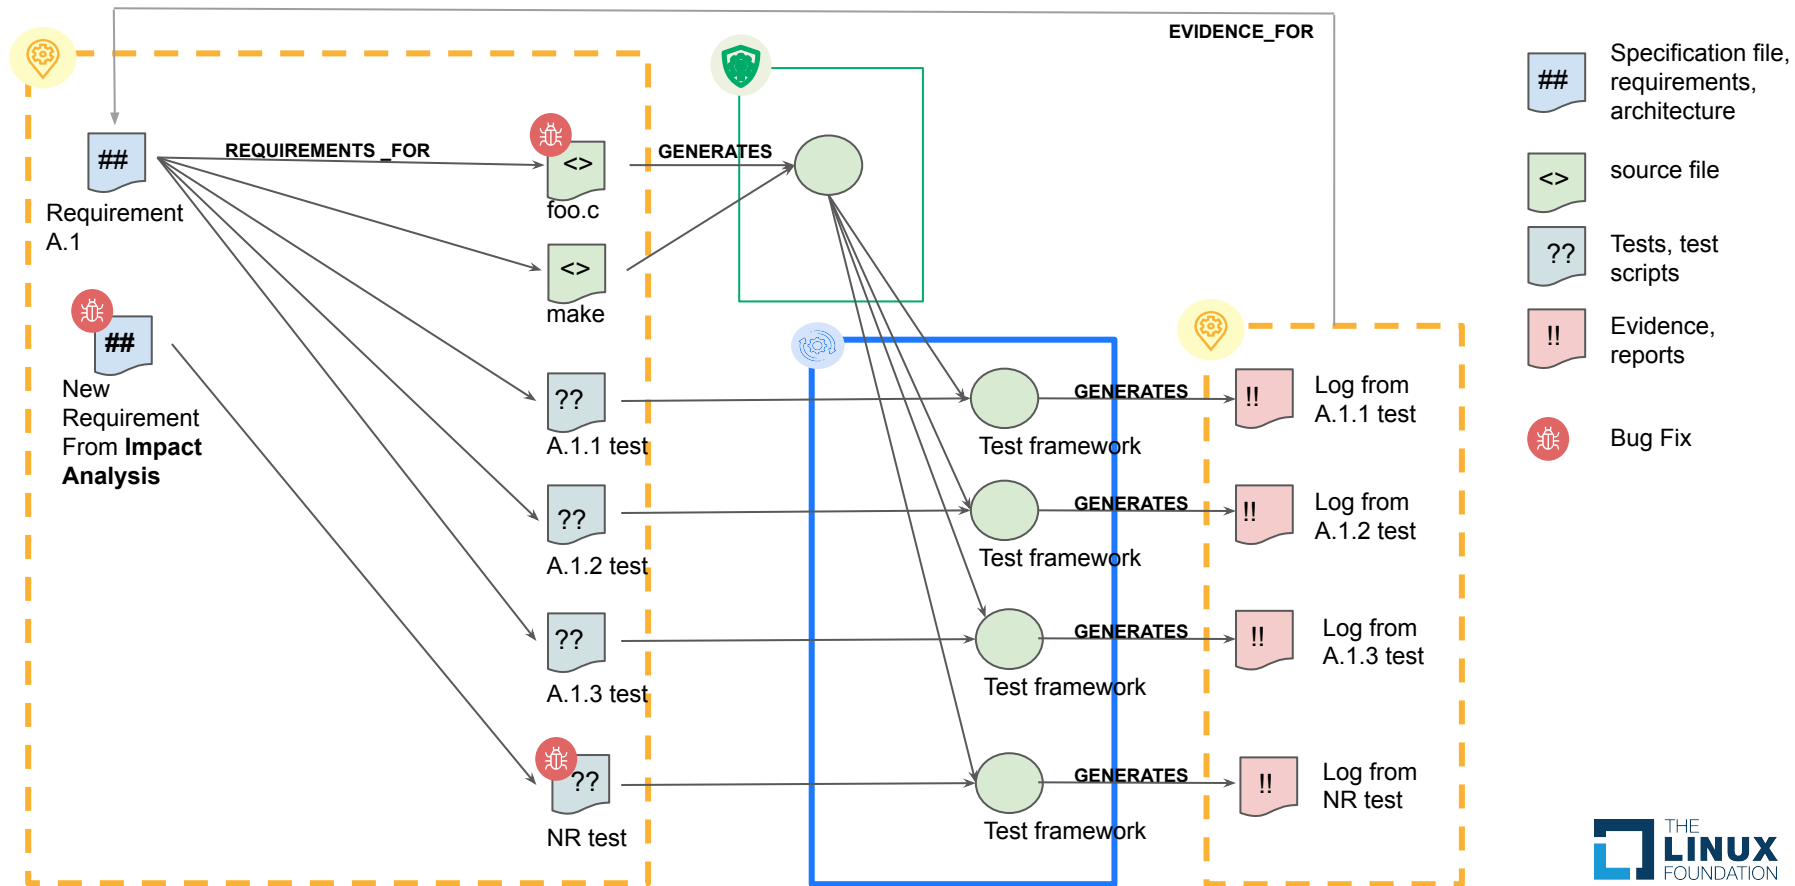
Requirement to Code to Tests to Evidence Traceability



Requirement to Code to Tests to Evidence Traceability



Requirement to Code to Tests to Evidence Traceability




Safety Profile Introduction (3.1 target)

- Also known as Functional Safety (or FuSa).
- Purpose is to link together all the safety artifacts (including code and relevant tests) with the aim of being able to automatically detect what a file update may need to force retesting.
- Goal is to support continuous certification of safety artifacts after security updates are applied.
- Overview can be found at: https://fosdem.org/2023/schedule/event/sbom_fusa/

For more information, contact: Nicole Pappler or Kate Stewart

Want to Help?



- If you have a **use-case** you want to make sure can be supported in the future SPDX specification,
 - join the SPDX tech team mailing list (<https://lists.spdx.org/g/Spdx-tech>),
 - open an issue in <https://github.com/spdx/spdx-spec> and
 - join in on the discussion!
- Try it, and let us know if you see issues.

Thank you! Questions?

How to Get Involved - PRs & Issues

<https://github.com/spdx>

- Specification
 - <https://github.com/spdx/spdx-spec> ← ISO submission format
 - <https://github.com/spdx/spdx-3-model> ← 3.0 development
 - <https://github.com/spdx/spdx-examples>
- Tooling
 - <https://github.com/spdx/tools-python>
 - <https://github.com/spdx/tools-golang>
 - <https://github.com/spdx/tools-java>
- License List
 - <https://github.com/spdx/license-list-XML>

Embedded Projects Generating SBOMs



Zephyr's west spdx

Presentation / Demo:

<https://www.youtube.com/watch?v=KYC3YpSu9zs>

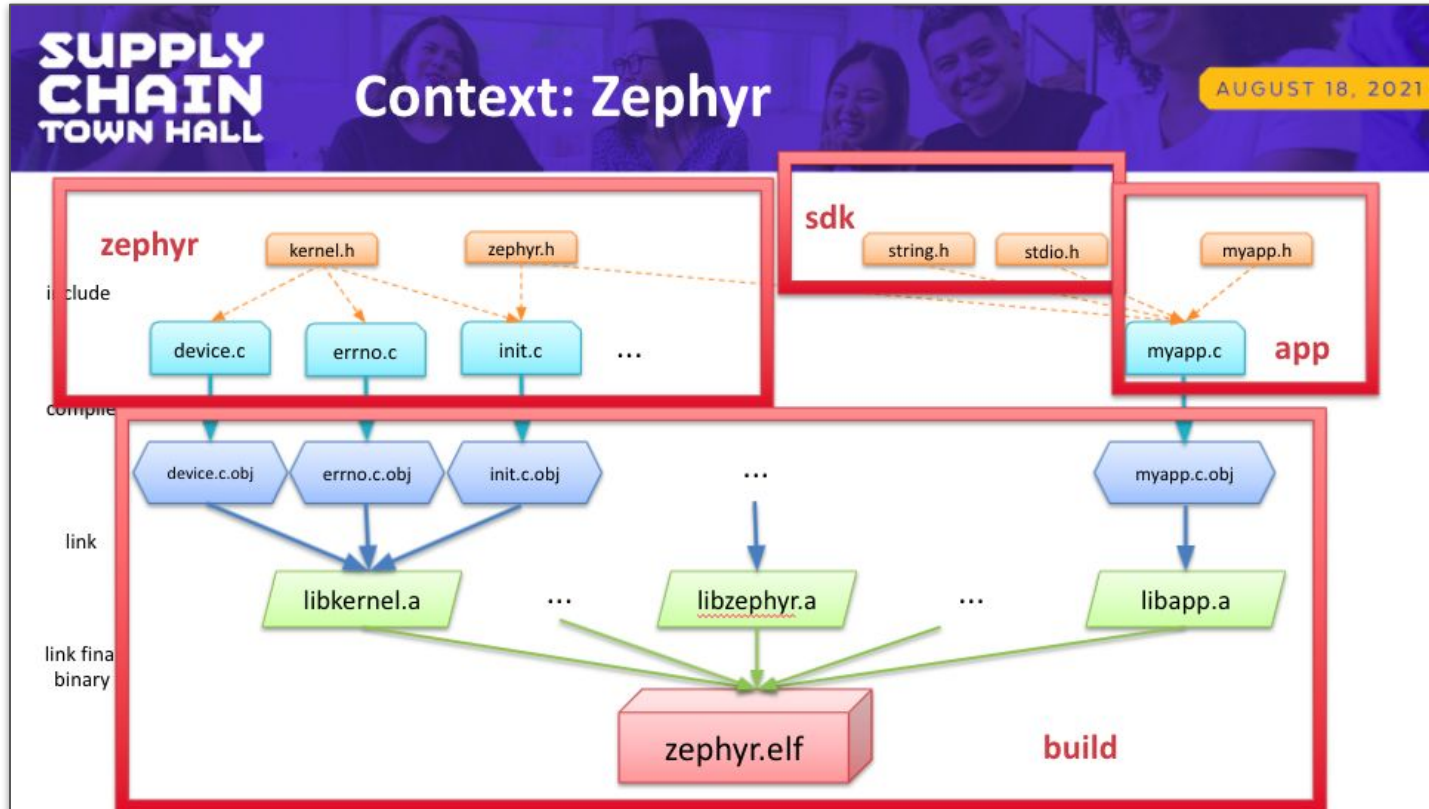


Yocto builds

Presentation / Demo:

<https://www.youtube.com/watch?v=y0N4FnkwTOY>

Relationship between SBOMs



SBOMs Included By Default ... Automatically

The screenshot displays the Renode Zephyr Dashboard. On the left, there is a sidebar with navigation options: STATUS MATRIX, BINARIES, ARCHITECTURE, BUILD DETAILS, and CONTACT US. The main area shows a list of boards under the heading 'BOARD NAME'. A search bar is at the top. Above the board list, there are four green status indicators: 64 PASSED, 99 PASSED, 105 PASSED, and 113 PASSED. A modal window is open, showing the SBOM for a TensorFlow package. The SBOM content is as follows:

```
TensorFlow
beaglev_starlight_jh7100-shell_module-build.spdx
SPDXVersion: SPDX-2.2
DataLicense: CC0-1.0
SPDXID: SPDXRef-DOCUMENT
DocumentName: build
DocumentNamespace: http://spdx.org/spdxdocs/zephyr-f51c77e5-c0b1-4354-b4fe-ce14d141768b/build
Creator: Tool: Zephyr SPDX builder
Created: 2022-01-18T12:43:08Z

ExternalDocumentRef: DocumentRef-app http://spdx.org/spdxdocs/zephyr-f51c77e5-c0b1-4354-b4fe-ce14d141768b/app SHA1: d64b4433e24afbe7c8f02bc6b0224b08ed9dc8e2
ExternalDocumentRef: DocumentRef-zephyr http://spdx.org/spdxdocs/zephyr-f51c77e5-c0b1-4354-b4fe-ce14d141768b/zephyr SHA1: fe54b93ae08866f4774f657aba960cbd4400a7a

Relationship: SPDXRef-DOCUMENT DESCRIBES SPDXRef-zephyr-final

#### Package: app

PackageName: app
SPDXID: SPDXRef-app
PackageDownloadLocation: NOASSERTION
PackageLicenseConcluded: Apache-2.0
PackageLicenseInfoFromFiles: Apache-2.0
PackageLicenseDeclared: NOASSERTION
PackageCopyrightText: NOASSERTION
FilesAnalyzed: true
PackageVerificationCode: fb5d511cc04828f9fb7ecc86074114eec9b651d

Relationship: SPDXRef-app HAS_PREREQUISITE SPDXRef-syscall-list-h-target
Relationship: SPDXRef-app HAS_PREREQUISITE SPDXRef-driver-validation-h-target
Relationship: SPDXRef-app HAS_PREREQUISITE SPDXRef-kobj-types-h-target
Relationship: SPDXRef-app HAS_PREREQUISITE SPDXRef-zephyr-generated-headers

FileName: ./app/libapp.a
SPDXID: SPDXRef-File-libapp.a
FileChecksum: SHA1: 5e8c866ca73e9c3dda61090d222834b30d484e6e
```

At the bottom of the modal, there are four buttons: NOT BUILT, NOT BUILT, BUILT, and BUILT.

Source: <https://www.linux.com/featured/enhancing-supply-chain-security-for-embedded-systems-renod-e-dashboard-for-zephyr-rtos-adds-new-software-bill-of-materials-sbom-capabilities-by-default/>