Anchore OSS Community Meeting

When: [Calendar Link]
12 PM ET / 16:00 UTC every other Thursday
Where: [Zoom Link] *(these meetings are recorded)*

Join Anchore OSS Community Google Group
To get write access to this document

Previous Recordings

---

1 August 2024

- [ ] Questions
    - Detecting vulnerabilities from an SBOM
    - CPE assignment for syft

18th July 2024

- [ ] CDN issues
    - Not blocked
- [ ] https://github.com/anchore/stereoscope/pull/268
- [ ] SBOM vs vulnerability scanner interoperability questions
    - Why do syft+grype work well together, triby-sbom-scanner+trivvy work well together, but not a mix in between? context : the interop point is an SPDX and CDX SBOM form the generator passed to the scanner
- [ ] Insert agenda items here!

4th July 2024

- No meeting today

20th June 2024

- [ ] Migration plan from Slack to Discourse
- [ ] PRs
    - https://github.com/anchore/syft/pull/2960

        - [ ] Brittle for making changes going forward to each user facing struct
        - [ ] Maintainer should come in and see if we can scope this to JUST relationships
            - [ ] Fix for the rust PR linked below is the main goal of this
            - [ ] Relationships were found to be unstable when sorting for Rust tests added below
    - https://github.com/anchore/syft/pull/2948
        - [ ] Config additions needed by Tools Team before merge (Maintainers are able to contribute this to make it easier)
        - [ ] Are we using the package verification field correctly?
        - [ ] Example Golang doing some of the things we would probably want to add:
            - [ ] Golang license download using the cache: https://github.com/anchore/syft/blob/main/syft/pkg/cataloger/golang/licenses.go#L127
            - [ ] Finding licenses instead of using a file.Resolver, but an `fs.FS`: https://github.com/anchore/syft/blob/main/syft/pkg/cataloger/golang/licenses.go#L139
            - [ ] Golang configuration is passed: https://github.com/anchore/syft/blob/main/syft/pkg/cataloger/golang/cataloger.go#L22
- [ ] Your item here?

6th June 2024

Agenda:

- [ ] Talk through cache busting options (or alternatives): https://github.com/anchore/syft/pull/2852
- [ ] (popey) Tyres on Community Discourse successfully kicked? Link/point to it from Slack/elsewhere? Currently pending ITHD-710 and ITHD-711 to move GH to Anchore account, and enable Google login. Might need to wait for these to be complete first?
    - 711 is done - so now you can login with github account
    - 710 is done - now you can login with google account \o/
- [ ] Your item here?

23rd May 2024

Agenda:

- [ ] (Alex) Should we consider https://github.com/anchore/syft/issues/661 completed? (reduce scope)
    - Probably not reduce the scope, but we can finish this later

- [ ] (Laurent) Curious about the lua PR
    - Looks like it's about to cross the finish line 🎉
- [ ] https://github.com/anchore/syft/issues/15
    - We have some thoughts on a direction here
        - [ ] File resolver that returns both the base and squashed representation, and a post processor to diff the results
        - [ ] A more holistic SBOM diff-patch approach for two appropriately scoped SBOMs (one that represents just the base, another that represents the whole artifact)

9 May 2024

Agenda:

- [ ] CPE confidence scores: https://github.com/anchore/syft/issues/2800

25 April 2024

Agenda:

- [ ] Plan for SPDX v3.0 adoption
    - Upstream changes needed in tools-golang: https://github.com/spdx/tools-golang/issues/237
- [ ] A few Syft questions:
    - CycloneDX transitive dependency relationships
        - [ ] https://github.com/anchore/syft/issues/572
        - [ ] https://github.com/anchore/syft/issues/2353
    - License information?
        - [ ] This would be great to include in the hash-based lookups we're thinking of
        - [ ] This is per-ecosystem [overarching issue]
- [ ] GrypeDB status update
    - NVD status – nvd overrides applied to OSS already: https://github.com/anchore/nvd-data-overrides
- [ ] As an OSS-project maintainer, what's an easy way to get started with SBOM?
    - sbom-action: https://github.com/anchore/sbom-action
- [ ] EPSS data in GrypeDB? [Keith to find or add ticket about this]

28 March 2024

Agenda:

- [x] ~~Chat through blocker for the PHP classifier PR https://github.com/anchore/syft/pull/2585~~
- [ ] Cataloging with regex:

- Keith's idea: having templating able to pull named regex parameters to put data in package structs
  - [ ] This would support things like separately finding major, minor, patch versions – such as redmine (a .rb file with separate major, minor patch)
  - [ ] This would also support non-binary things, such as joomla php files (composer finds dependencies, but not the overarching software)
  - [ ] Testing right now has scripts in only binary cataloger, would need to be made generic/moved to testutils somehow
- [x] ~~Ready for review: https://github.com/anchore/syft/pull/2613~~
  - This gets 60% or so of rockspec, but rockspec is in lua so there will be a follow-on to improve this

14 March 2024

Agenda:

- [ ] GrypeDB – discuss plans to fill the gap when NVD isn't providing more structured CVE data
  - Blog: https://anchore.com/blog/national-vulnerability-database-opaque-changes-and-unanswered-questions/
  - GitHub repo https://github.com/anchore/nvd-data-overrides/
    - [ ] This is meant to be a temporary option to fill the current NVD gap - Just includes the configurations override
      - [ ] Be very very clear that this is a short term thing
    - [ ] We need a new repo for a more robust option (that contains more than just NVD configurations)
      - [ ]
  - What sort of override format would we prefer?
    - [ ] PURL?
      - [ ] There are things that currently fall outside the PURL spec
    - [ ] The format used to annotate the data doesn't have to be the output format
      - [ ] For example we could have a format that isn't OSV or CVE5, but could output those formats
    - [ ] Do we want to import all the vulnerability data, or do we only want to capture certain extra data (overrides)
    - [ ] We could wait until after vulncon to decide how to generate overrides
    - [ ]
  - How can the community add to the dataset?
    - [ ] We need to create a process
  - Using LLM/ML to classify CVE descriptions into package details
    - [ ] Anchore has an experiment for this

        - [ ] The intent of this process isn't to automatically add data, but to let a human work a lot faster (instead of having to figure out what the cpe/package is, we get a head start from the AI)
        - [ ] Making the review process as simple as possible is needed
        - [ ]
    -
- [ ] Q&A
    - Your questions here
- [ ] Your topic here

29 February 2024 Special leap-year edition!

Agenda:

- [ ] Q&A
    - Your questions here
- [ ] Your topic here

15 February 2024

Agenda:

- [ ] Demo of grant, Anchore's new open source tool for license analysis.
- [ ] **Q and A**
    - **No Q and A today -**
- [ ] **We'll release a Youtube Video of the new tool for others to consume!**
- [ ] *Your topic here!*

1 February 2024

Agenda:

- [ ] **Q and A**
- [ ] Vunnel data and vulnerability data
    - CPE matching / When we use the NVD namespace
    - We stopped searching for vulnerabilities by CPE (common platform enumerator) for ecosystems covered by the Github Security Advisory Database because their DB is indexed by ecosystem. This really helps prevent overlaps and false positives when there are common names across ecosystems. Example: a docker client in Node and a docker client in Rust might have the same name so their vulnerabilities would show up against each other.
        - [ ] Screen shot showing the ecosystems supported by GHSA

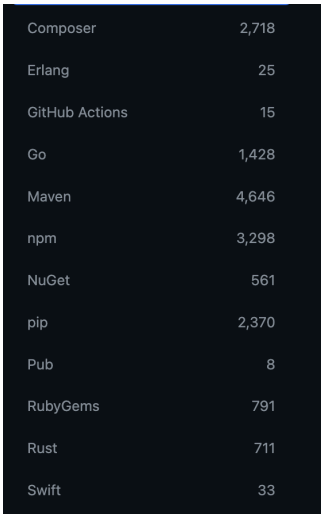

    - Direct vs Indirect Matching in grype
        - [ ] Matching with a name found in the SBOM vs matching against a name NOT found in the SBOM. For operating systems you can file vulnerabilities under the SOURCE package and not the package you're looking at. When something is indirect we searched by package name (dependency) and source name (parent package) to try and find a vulnerability. An example of this would be something openssl and a dependency used to build it. If we find the dependency then sometimes there are cases where a vulnerability in that dependency is filed against openssl. Grype will search by both names. If a vulnerability is found by searching the parent (openssl) instead of the original dependency in the SBOM we call it an indirect match
    - Grype is still best when using `syft -o json`
    - Discussion on including symbol tables in SBOM and how they get a little to large
        - [ ] This was referenced https://pkg.go.dev/golang.org/x/vuln/cmd/govulncheck
    - Follow up to the question about SBOMs that contain package names referencing a sha:256
        - [ ] https://anchorecommunity.slack.com/archives/C027JE5M345/p1706810402133869
- [ ] https://github.com/anchore/vunnel/pull/473
- [ ] Potential bug in matching against java when no PURL available Josh Bressers
- [ ] **New API Examples:** https://github.com/anchore/syft/pull/2517

18 January 2024

Agenda:

- [ ] Talking about plugfest

- Instructions: https://docs.google.com/document/d/1USn3XJ25d6l8nMqBTnsqoovJIW36KupODiqvo8ZpqPo/edit
- Submissions: https://drive.google.com/drive/u/0/folders/1Ujxp8w7dhrL6TNj5NxcaASbqPSoVTP1Y
- Analysis: https://drive.google.com/drive/u/0/folders/1IWMHHTCtHIFDUpgJM8gpaoratalMoH8l

4 January 2024

Agenda:

- [ ] Happy new year!
- [ ] https://github.com/anchore/syft/pull/2469
  - Take a look at https://github.com/anchore/syft/pull/2444
  -

21 December 2023

Agenda:

- [ ] https://github.com/anchore/syft/releases/tag/v0.99.0
- [ ] We talked about the state of licenses in the two formats as

7 December 2023

Agenda:

- [ ] Syft convert is inheriting a tool version but not name (in the descriptor section)
- [ ] SPDX and CycloneDX output is not writing relationships in equivalent ways (specifically pkg-to-pkg relationships)
- [ ] Capturing dependencies: https://github.com/anchore/syft/issues/572
- [ ]

25 November 2023

Canceled* -- *Happy Thanksgiving 🦃

9 November 2023

Agenda:

- [ ] Any update on https://github.com/anchore/syft/issues/1711 ?

26 October 2023

Agenda:

- [ ] --file vs scheme vs URI input https://github.com/anchore/syft/issues/1783
    - `syft registry:myimage:latest`
        - [ ] Hold over behavior
    - `syft --from registry myimage:latest +1`
        - [ ] Anchorectl is already doing this
        - [ ] Is separable
    - `syft registry://myimage:latest`
        - [ ] Might overlap with future input that would start with https:// (or other URIs)

- [ ] Should we replace the “packages” command with another noun or verb? https://github.com/anchore/syft/issues/516#issuecomment-996999477
    - Current thoughts are that syft is centric to SBOMs in general and sticking to verbs on the CLI is the next direction. This implies “create” (an in create an SBOM)
        - [ ] +1 on verbs
            - [ ] Create
                - [ ] Create doesn’t operate on the argument (english-wise?)
            - [ ] Scan +3
    - `syft packages myimage:latest`
        - [ ] Hold over behavior
    - `syft myimage:latest (root is aliased to the packages command)`
        - [ ] `+2 on keep the alias`
    - `Syft convert sbom1.json -o spdx > sbom2.json  (verb example)`
    - `Syft create myimage:latest`

- [ ] False negatives and our current approach False negatives when scanning debian trixie/sid images from Dockerhub · Issue #1446 · anchore/grype · GitHub
    - If we keep the mappings we should try to do so in grype-db and not grype, so that on each codename introduced we don’t need a grype release
    - Syft:
        - [ ] Could we additionally update syft to account for this?
        - [ ] Maybe we could fix this in syft
        - [ ] Implies we are shipping a data side car regularly (which contains the codename mapping)
        - [ ] Can make the mapping configuration driven in grype (or syft)

        - [ ] **Without the distro version in syft all of the purls don't include version information**
        - [ ] How often are the codenames used (once a release? 6 months? Once a year)
        - [ ] **We could simply add the hardcoded mapping in syft to resolve this in the short term**
        - [ ] **Can additionally add this and grype and syft (so that grype doesn't depend on this replacement being done upstream on it's behalf)**
        - [ ] **If we do this in syft, there should be a test that checks a remote endpoint to ensure that mapping is up to date**
    - Short term: do the bold steps?
        - [ ] They are quick and get the job done, however they are not data driven
    - Long term: make this data driven (grype-db or syft data side car)

```
❯ syft -q -o json debian:trixie | jq .distro
{
  "prettyName": "Debian GNU/Linux trixie/sid",
  "name": "Debian GNU/Linux",
  "id": "debian",
  "versionCodename": "trixie",
  "homeURL": "https://www.debian.org/",
  "supportURL": "https://www.debian.org/support",
  "bugReportURL": "https://bugs.debian.org/"
}
```

- [ ] Chicken Egg (Labels and Quality Gate)

12 October 2023

Agenda:

- [ ] A discussion of our decision to stop matching on CPEs for most vulnerability matches (https://github.com/anchore/grype/pull/1412) in an upcoming release of Grype.
    - Released! https://github.com/anchore/grype/releases/tag/v0.71.0
- [x] ~~Syft 1.0 breaking changes to the json format: curious to hear folks opinions on these changes~~
    - https://github.com/anchore/syft/issues?q=is%3Aopen+is%3Aissue+label%3Ajson+milestone%3A%22Syft+1.0%22 .. in particular:
        - [ ] https://github.com/anchore/syft/issues/1842
        - [ ] https://github.com/anchore/syft/issues/1835
    - Feedback

        - [ ] Maybe defend what features belong under the config flag
        - [ ] Maybe have the syft config flag be specific to the schema version instead of just “legacy”
- [x] ~~Updates on https://github.com/anchore/grype/issues/1441~~
    - We might not need to do this in grype as this isn’t useful to end users, in vunnel or grype-db makes more sense
        - [ ] https://github.com/anchore/grype-db/blob/ef7ddfd459086a46eec2f69856c9bdd80eb3d2d8/pkg/process/build.go#L152-L159 …. Alex Goodman post a better answer on the issue
- [x] ~~Handling ubuntu vuln db out of support~~
- [ ] *Your topic here!*

28 September 2023

Agenda:

- [ ] Discuss configuration of vulnerability severity display when different data sources have different severities: https://github.com/anchore/grype/issues/1378
    - This is another signal that related vulnerabilities section is off (we only show upstream, but we could show additionally downstream too)
    - This applies to what formats/feature?
        - [ ] Table
        - [ ] --fail-on
        - [ ] JSON?
    - Really we want to allow configurability to show the single severity for possibly more than one record
- [ ] Signing artifacts https://github.com/anchore/grype/issues/1513
    - Start with checksum signing, stretch goal for docker images + sbom
- [ ] Base image identification
    - https://github.com/anchore/syft/issues/1199
    - Suggestion: allow list of known base images + digest for popular images
        - [ ] Publish on Rekor https://github.com/anchore/syft/issues/1199#issuecomment-1271147467
            - [ ] Known issue: anyone can write, you need a trusted list of publishers
        - [ ] Option: should docker build include this?
    - Investigate:
        - [ ] How does the ADD file:... id get generated?
        - [ ] Rekor: what do we need in syft to use Rekor in a safe fashion (I think a list of emails/domains)
        - [ ] Does buildx have an issue to track metadata for this?

- [ ] Are there any conventions around the LABEL line in a dockerfile and how we can track this info post-build?
    - [ ] https://dockerlabs.collabnix.com/beginners/dockerfile/Label_instruction.html
- [ ] Are there any conventions around annotations that we could leverage? This would be verifiable if online
    - [ ] https://github.com/docker/roadmap/issues/243

- Actions:
    - [ ] Does stereoscope have the ability to persist docker label information for syft to use?

14 September 2023

Agenda:

- [x] ~~Questions about related vulnerabilities and how they are reported.~~

31 August 2023

Agenda:

- [ ] Grype: add configuration of severity display
  https://github.com/anchore/grype/issues/1378
    - [ ] Add configuration to prefer severity from various data sources
        - [ ] Possible values (probably mutually exclusive):
            - [ ] Prefer highest severity
                - [ ] Ideally this is when there are multiple IDs in a source (like a distro advisory). In the case of differing severities for one ID, defaulting to one trusted source could be preferred
            - [ ] Order source categories (e.g. trust vendor first, other advisory sources next [ghsa], nvd last)
            - [ ] “Don’t care about any severities that require physical access”
        - [ ] When we allow adding severity sources into grype db
            - [ ] Should we try to tackle this first before adding configurable severity?
        - [ ] The table output has one severity
            - [ ] Should this change? Should that be configurable?
            - [ ] Maybe a little more info in the default output
            - [ ] This gets more interesting with the grype explain json output (and potentially future CVE oriented output): this turns into a rollup problem, how do we be transparent about where each severity came from.

- [ ] Migrating the ecosystem to go 1.21 (go toolchain directive) https://github.com/anchore/syft/issues/2066

17 August 2023

Agenda:

- [x] ~~Discuss "grype explain", a proposed command to show more information about a specific vulnerability match and possibly how to remediate: https://github.com/anchore/grype/issues/1342~~ ~~(Will M.)~~
    - [ ] Package information could be higher (highest)?
    - [ ] OS package relationship info to eliminate packages
- [x] ~~Protobuff for grype~~
    - [ ] Json schema for both input and output
- [x] ~~CVSS v4~~ ~~https://www.first.org/cvss/v4-0/~~
    - [ ] Add validation around the CVSS … preferably in the grype output
- [ ] Discuss the idea of having Grype report "negative matches": package/vulnerability pairs where grype found evidence that the package is **not** affected by the vulnerability. https://github.com/anchore/grype/issues/1426 (Will M.)
- [x] ~~Take questions about the upcoming package metadata breaking changes https://github.com/anchore/syft/pull/1983~~

3 August 2023

Agenda:

- [ ] Breaking changes to package metadata are inbound https://github.com/anchore/syft/pull/1983
- [ ] SBOM analysis: https://sbom.seclab.cs.hm.edu/#/sbom-files
    - [ ] SBOM convert
        - [ ] Problem: some SBOMs don't report the format version

18 July 2023

Canceled

6 July 2023

Agenda:

- [ ] https://github.com/anchore/syft/issues/1896 – cataloging Github workflows
- [ ] Prioritizing OSS community issues in the backlog column

22 June 2023

Agenda:

- [ ] Upcoming Syft 1.0 changes / discussion https://github.com/anchore/syft/milestone/2
    - [ ] Guarantee package ID stability between syft versions
        - [ ] This means we'd need to ignore new fields
        - [ ] What about semantic names to fields?
        - [ ] Alex Goodman I think zach needed this at some point (check this)... we should capture these use cases
        - [ ] To make this easier we could change the hash function to an allow list (today it's a deny list based on struct tags)
    - [ ] Non-breaking JSON schema changes
        - [ ] Should we remove guarantees of the pkg.Metadata?
            - [ ] Using grype here as an analog, there are only a few fields used
        - [ ] Another way to approach this is duplicate fields in the short term
    - [ ] Maybe there is a connection between the JSON schema fields and ID stability
    - [ ] Syft versioning*
- [x] ~~Data race in Stereoscope: https://github.com/anchore/stereoscope/issues/184~~
- [x] ~~What is the connection between vulnerabilities and related vulnerabilities? Sometimes we see several related vulnerabilities corresponding to one vulnerability, like "ELSA-2022-9221" with "CVE-2021-20232", "CVE-2021-3580", and "CVE-2021-20231". Is it the same vulnerability under a different name?~~
    - [ ] --by-cve

8 June 2023

Agenda:

- [ ] Questions about Syft backward compatibility/interface stability before 1.0 as it relates to Tekton plugin (requested by Al H.)
- [ ] Conda discussion: https://github.com/anchore/syft/issues/932
    - [ ] TODO: Reach out to Jerimiah in community slack to help with pairing/starter efforts
- [ ] Grype Data Fields
    - [ ] Set an expiration date on the ignore field

        - [ ] https://github.com/anchore/grype/blob/main/grype/presenter/models/ignore.go#L10-L14
        - [ ] Would it be possible to add a field here that sets an expiration date
        - [ ] Confusion between `Exclusions` and `Ignore Rule`
        - [ ] https://github.com/anchore/grype/issues/452
        - [ ] User Notes (greg):
            - [ ] For each entry - a comment has a hash (could be anything)
            - [ ] Day Waivered entry per ignore is suggested
            - [ ] Global counts down and entries show up dynamically based on the way waivered entry
            - [ ] Reason to ignore (comments)
- [ ] More signals for grype explain:
    - [ ] Show specific package details for a given vulnerability entry on the CLI
- [ ] Frustrating to re-run grype over and over
    - [ ] How do we make it clearer to users that if they run grype with an sbom it's faster
    - [ ] Should we cache the original sbom?

25 May 2023

Agenda:

- [ ] Upcoming JSON schema changes (will be done incrementally before syft 1.0)
    - [ ] Inconsistent MetadataType names https://github.com/anchore/syft/issues/1844
    - [ ] Rename location fields https://github.com/anchore/syft/issues/1835
    - [ ] Rename top level json fields https://github.com/anchore/syft/issues/1842
    - [ ] Should breaking json schema changes require a breaking change to syft?
        - [ ] So far, we're assuming "yes" moving forward beyond syft 1.0
        - [ ] Maybe "no" if we guarantee backwards compatibility in decoding
        - [ ] We could add multi-version encoding support, default encoding to locked major versions, but allow for "syft -o json@dev" (e.g.) to express breaking functionality in an opt-in way.
- [ ] Opportunity to help shape ELF metadata for non-OS-packaging cases https://github.com/anchore/syft/issues/1821
- [ ]

11 May 2023

Agenda:

- [ ] Should we move forward with https://github.com/anchore/syft/issues/572 without --with-tools support (the original plan)?
    - [ ] Question: reconcile with-tools and without-tools

        - [ ] These mechanisms feel mutually exclusive… (we raise up relationships differently for each method)
    - [ ] Probably move forward with the local-only approach, this can be independent work relative to --with-tools

- [ ] Capturing Known-unknowns https://github.com/anchore/syft/issues/518 , are there any use cases we should focus around with this feature
    - [ ] Minimal guidance from https://www.ntia.doc.gov/files/ntia/publications/sbom_minimum_elements_report.pdf
    - [ ] Suggestions / Examples:
        - [ ] Binary cataloger, can't find a versions (e.g. but we know it's NGINX)
            - [ ] "Partial packages" that we drop during cataloging but could capture as a known-unknown ← there are a few places similar to this across catalogers
            - [ ] Go packages where we have no version info for the main module (see also https://github.com/anchore/grype/blob/d74e85385ca0e83d28e14d2a6851c21cdf805638/grype/matcher/golang/matcher.go#L48-L50)
        - [ ] "We think this is an alpine image but there is no alpine DB"
        - [ ] We think this is a python site-packages dir, but there might be missing egg metadata
            - [ ] Given dir structure, there should have been known metadata that is gone
        - [ ] We found a dir with source code but no ecosystem package manager metadata (.py without requirements.txt for instance)
        - [ ] Expectation statement (user input) vs what was found (sbom)
        - [ ] List files that are not claimed by a package
    - [ ] Probably put this into a separate section so that it's easily separable
    - [ ] Should be configurable such that it can easily be turned off in case it's too verbose
        - [ ] Maybe we have this off conditionally based on the source type? (e.g. dir scans)

- [ ] Capture distro information in grype matchers to evaluate package qualifiers

27 Apr 2023

Agenda:

- [ ] Discussion about the refactoring the source package https://github.com/anchore/syft/issues/558#issuecomment-1525832828

    - [ ] Lets go forward with a draft PR concept
- [ ] Polling interest in FIPS compliant builds of Syft and Grype (enforcement of cryptographic methods used)... is there interest?
    - [ ] What if we only released a container image build? Is there desire for static binary release (no investigation has been done on the LoE for this)?
    - [ ] Using UBI image to release tooling
        - [ ] https://developers.redhat.com/articles/2022/05/31/your-go-application-fips-compliant#
        - [ ] https://gitlab.com/gitlab-org/security-products/analyzers/secrets/-/blob/master/Dockerfile.fips
    - [ ] Alex Goodman make an issue for this
- [ ] [help wanted] Interpretation of Portage license handling fixes… can we translate these expressions to SPDX license expressions?
    - [ ] https://github.com/anchore/syft/pull/1763/files#diff-bf3cacb30ec9707448aeb6b23b0d80b6ae8cd493194e415cc97ac677124c5188R47
    - [ ] https://github.com/anchore/syft/pull/1763/files#diff-bf3cacb30ec9707448aeb6b23b0d80b6ae8cd493194e415cc97ac677124c5188R40
    - [ ] https://github.com/anchore/syft/pull/1763/files#diff-bf3cacb30ec9707448aeb6b23b0d80b6ae8cd493194e415cc97ac677124c5188R53
    - [ ] Path forward:
        - [ ] Probably just the mandatory licenses
        - [ ] In this case we should keep the license string in the metadata struct
        - [ ] Capture this as an issue and reach out to gentoo?

---

13 Apr 2023

Agenda:

- [x] ~~Chit chat about vunnel how do vulnerability databases work? (Henry)~~
- [x] ~~Embed fs.FileInfo in Metadata struct? https://github.com/anchore/stereoscope/issues/171 (Ad)~~
- [x] ~~Build tooling query ~https://github.com/anchore/syft/issues/1736~ https://github.com/anchore/syft/issues/1562~~
- [x] ~~Purls! Upstream source packages….~~
    - [ ] Should we get this in to the purl spec?
    - [ ] Should we add multiple purls to the SBOM?

---

30 Mar 2023

Agenda:

- [x] ~~Evidence Relationships https://github.com/anchore/syft/pull/1698~~
    - [ ] How do other tools express files evident by?
- [x] ~~Config option for image source~~
    - [ ] Consider --from like anchorectl
    - [ ] Nit: image-pull-source -> default-image-pull-source

---

16 Mar 2023

Agenda:

- [x] ~~Henry would like to talk about Gradle Dependency scanning: Draft: feat: initial gradle implementation by henrysachs · Pull Request #1407 · anchore/syft (github.com)~~
    - [ ] Seeing “Finished generating 'example-java-app' in 1m 20s.” with a lot of output. It looks like there may be a cache miss (erroneous?) that should be accounted for.
- [x] ~~Talk about local (non external license files in golang - https://github.com/anchore/syft/pull/1645~~)
    - [ ] Capture configuration refactors into a separate issue
    - [ ] Lets try to add explicit blockers on the PR (or requests for change / expectations)
- [x] ~~Final update being pushed here https://github.com/anchore/syft/pull/1554~~
- [x] ~~Seeing a null severity for CVE-2023-28154~~

---

2 Mar 2023

Agenda:

- [ ] Avi’s PRs: https://github.com/anchore/syft/pulls/deitch

---

16 Feb 2023

Agenda:

- [x] ~~Syft License Revamp~~
    - [ ] Google License classifier

- [ ] https://github.com/google/licenseclassifier/blob/bb04aff29e72e636ba260ec61150c6e15f111d7e/README.md
- [ ] https://github.com/anchore/syft/issues/1577
- [ ] https://github.com/anchore/syft/pull/1554
  - [ ] Other License Field
    - [ ] Include all the text of the thing that is non concluded. This is the actual text of the license set as the field value in the SBOM
    - [ ] Ask the author to give the SPDX file - machine readable makes it as the one authoritative source
    - [ ] Human reading does not rely on tool when using the SPDX ID. Concluded was intended for humans to go back and say yes this is the license and fill in this
- [ ] Abstract syntax tree of the identifier:
- [ ] Nodes are license ID OR a custom License ID
- [ ] The path between them are the operators AND, OR or WITH
- [ ] 3.0 of SPDX where there is one and only one way to specify something in the document
- [ ] 2.3 License information in file field
  - [ ] I found this but don't know what it means
  - [ ] For 3.0 there will be an equivalent for found in file license information
- [ ] Canonical SPDX representation
- [ ] Multiple different text representations for the same tree structure
- [ ] []string{"expression1", "expression2"}
  - [ ] Some ecosystems that in package metadata declare a license list see NPM
  - [ ] Initial cut of looking at expressions was deferred to splitting on AND
  - [ ] **License concluded is where the creator of the SPDX document decides that some part of the expression is redundant. Example (MIT OR APACHE) can be simplified to MIT**
  - [ ] **Concluded is a human focused field and not the output of a pur SCA tooling**
  - [ ] GPL and LGPL - given the ability to simplify down to GPL when you are consuming this package on issuing the package
  - [ ] Not using concluded unless the the packages are actually bundled
  - [ ] Strings for representing expressions will not be deprecated
- [ ] Grype source scanning .NET projects
  - [ ] https://github.com/anchore/syft/issues/1522 similar to https://github.com/anchore/syft/issues/690
    - [ ] Use Nuget from .csproj files, maybe?

---

2 Feb 2023

Agenda:

- [x] ~~Revisit the custom attestation type: https://github.com/anchore/syft/issues/1532~~
  - [x] ~~We should check if cosign is working now for non-custom types~~
    - [ ] Only works if explicitly provided
  - [x] ~~Add syft attest --type ?~~
    - [ ] Merged PR from 1532 and addingtype flag
  - [x] ~~Either way, syft probably shouldn't force an opinion here for the type~~
- [x] ~~SPDX licenses that are not in the official license list: https://github.com/anchore/syft/issues/1529~~
  - [ ] Start populating the other-licenses section (for non exact matching license info). We can reference these within "license declared/concluded" (point to the other section).
- [x] ~~Preview large changes from~~
  - [x] ~~https://github.com/anchore/stereoscope/pull/154~~
  - [x] ~~https://github.com/anchore/syft/pull/1510~~
- [x] ~~Output specific versions of formats (e.g. SPDX 2.2, 2.3)~~

---

19 January 2023

Agenda:

- [x] ~~Resolver path filters:~~
  - [ ] https://github.com/anchore/stereoscope/pull/151
  - [ ] https://github.com/anchore/syft/pull/1447
- [x] ~~Surfacing symlinks as owned files instead of resolved absolute files [#1387]~~
  - [ ] We might have another path forward by adding additional coordinates in a separate location attribute

---

5 January 2023

Agenda:

- [x] ~~SPDX DocumentNamespace: https://github.com/anchore/syft/issues/1421~~
    - Can use .invalid URI to ensure it's never valid / intentionally host a landing page that explains
    - We could allow for a user override for this field, either the entire value or a prefix, to allow for hosting
    - Could the unique portion be a UUID or a digest?
        - We could have the digest of the SBOM object in memory
- [ ] SBOM builder configuration is going to be merged soon, looking for feedback
    - https://github.com/anchore/syft/pull/1383
    - https://github.com/anchore/syft/issues/558
    - Could we provide a useful default for getting a logger
        - We should also setup stereoscope nicely for folks as well

---

22 December 2022

Agenda:

- [x] ~~Injecting cataloger configs in https://github.com/anchore/syft/pull/1383~~
    - Maybe generics or another interface for configs… meet up later
- [ ] Overlapping package types: https://github.com/anchore/grype/issues/1044
    - Extract type via purl
    - Investigate a matcher config that allows users to augment based on package and relationship triggers
    - We should add more warnings about possible data problems
        - Maybe there should be a more up-front way to present these instead of just the log
        - Help out in an audit trail to show when decisions are made
        - Alex Goodman create an issue
- [ ] Grype python version comparisons: (Grype 1034)

---

8 December 2022

Agenda:

- [x] ~~Surfacing symlinks as owned files instead of resolved absolute files [#1387]~~
    - Options:
    - One solution is to only surface the symlink files
        - We lose information about the absolute file that it points to. We are correct in that there is no ownership overlap
    - Add more metadata into the location ← **This one**
        - Full symlink resolution by enhancment
            - sa => sb => absolute path
            - Ownership is for sa
            - But package knows how to traverse to the absolute path from the owned files
            - SPDX path sa DYNAMIC_LINK => sb =>DYNAMIC_LINK => absolute path
- [ ] Usecase: Where in the dockerfile did your vulnerable package come from
    - Owned files that are just indirect references of each other make this question harder
    - One owned file that can do the traversal makes it easier
- [ ] This is how buildkit does it
    - It never sees multiple layers (this is logistically insane)
    - Gives syft the root file system and goes and scans this (no layers)
    - Build kit goes through this scanned system and adds layer information after. This process makes the multiple ownership model break
- [x] ~~Question about how Grype builds its vulnerability database from the vulnerability datasources~~
    - Which grype database vulnerability namespaces have vulnerabilities with fix versions? (eg in the database for *nvd:cpe* namespace I only saw null *fixed_in_versions)*
    - Which grype database vulnerability namespaces have cvss? (eg in grype database for *github* namespaces they do not have entries in the CVSS array 904)

---

24 November 2022

…Canceled! Have a happy Thanksgiving!

---

10 November 2022

Agenda:

- [ ] Grype Image including database built nightly
    - Add a ticket for this, see Chris' comment here: https://github.com/anchore/grype/issues/837#issuecomment-1205509596
        - https://github.com/anchore/grype/issues/987
- [ ] SBOM cataloger - https://github.com/anchore/syft/pull/1029
    - Docker official images have a need to put an SBOM for things that are not being found by scanners, want to put SBOMs in the filesystem
    - Additive is probably fine for what Docker wants to do
    - Moving this forward: see if there is anything Anchore can do to clean this PR up and get it to the finish line
    - File ownership information is very important
- [ ] https://github.com/anchore/syft/pull/1269 (Add --name option to override name in output #1269)
    - Keith have a look at this PR
        - Approved!
- [ ] https://github.com/anchore/syft/issues/1129 – Maven variable handling
    - Need to determine if it's possible to do network requests rather than just static analysis
    - Might be able to introduce a feature flag to do this
- [ ] https://github.com/anchore/syft/issues/1115 – Enriching SBOM data from external sources
    - Need to have some core developer conversations about the direction of Syft
- [ ] False positives
    - Postgres with blank version

Notes:

-

---

27 October 2022

- Reporting Relevant CVE number for GH https://github.com/anchore/grype/issues/204
    - This specific issue is not a problem, as the GHSA was updated with CVE information
- Update on Grype database missing CVSS vectors for GHSA https://github.com/anchore/grype/issues/904
- Followup for almost duplicate packages being found by catalogers with internal overlap https://github.com/anchore/syft/issues/1162

    - The biggest pain here isn't in Syft but in Grype, having multiple of "the same" vulnerability. Maybe we can modify Grype instead to merge the same vulnerability and have multiple locations under a single element – looking at the Grype JSON output, we (except for JavaMetadata virtualPath, we could consolidate all these and just include all the locations together)
- Question about false positive with some npm packages https://github.com/anchore/grype/issues/890
    - Maybe add CPE exclusion flag to be able to exclude things like ***:chainsaw**

---

13 October 2022

- [x] ~~Handling **alpine:edge** and similar in Grype matching (from this community discussion)~~
    - [ ] Fix warnings such that there is not duplicate distro warnings
    - [ ] V3 alpine will have per-package distro info
    - [ ] Namespace logic should account for _alpha/_beta mapping to edge
        - [ ] Confirm if we need to only do this mapping for the latest version? Or do we need to look at the DB for all distinct namespaces?
    - [ ] https://github.com/anchore/grype/issues/954
        - [ ] No CPE in the index yet
        - [ ] Trivy is finding this, they are using secdb as an advisory feed. (an old grype bug)
- [ ] Additional APK data in SPDX https://github.com/anchore/syft/issues/1250
- [ ] Collecting comments about open sourcing the data pipeline for grype:
    - [ ] If you wanted to add another vulnerability data source, what language would you prefer to write the parser in?
        - [ ] It's less about the language in question and more about if we can moderate changes to the core grype data pipeline. That is, we should be using external source code that we can't directly change. But we also don't need to require that all surface areas of data entry need to be controlled for all grype users (if you are building your own custom db with your own sources for instance=)
    - [ ] Deprecation approach to grype DB schemas. Context: we build and publish a new vulnerability database for all DB schemas that grype has ever processed (today there are 4). As we add more, this will become more expensive to create and store. In order to make this more scalable we want to start deprecating DB schemas (which implies that we would no longer be supporting grype versions of those schemas).
        - [ ] Always keep the latest version of a deprecated schema (never deleted)
        - [ ] Deprecate data updates to schemas, not the schemas themselves

- [ ] A great trigger point for more aggressive deprecation behavior is when we open source the data pipeline
- [ ] Add more banner warnings into grype about upcoming deprecations to schemas

- [ ] [not covered last time?] Docker daemon provider arch variant https://github.com/anchore/stereoscope/issues/143
  - [ ] M1 mac user to validate when we fix the empty string for variant in stereoscope

---

29 September 2022

- [ ] Chapman picking up https://github.com/anchore/syft/issues/1162
- [x] ~~Linked SBOM discovery via rekor~~ ~~https://github.com/anchore/syft/issues/1159~~
  - [ ] How do we reconcile two identical rekor entries that advertise 1st vs 3rd party assertions?
  - [ ] Same question where 1st party uploads twice but the fresher is worse
    - [ ] Both cases of fresher broken vs fresher worse (describing the same software)
    - [ ] ABI fingerprinting
    - [ ] Issue: Lots of build processes will patch shared libraries so hashes are no always ideal
    - [ ] Performance research on Rekor batch queries
    - [ ] External reference dead data of dropping a link
    - [ ] `Relationship: SPDXRef-DOCUMENT AMENDS DocumentRef-SPDXA:SPDXRef-DOCUMENT Is this what we want for the rekor cataloger?`
- [x] ~~Exclude packages option~~ ~~https://github.com/anchore/syft/issues/1229~~
  - [ ] Exclude a file as well came up
    - [ ] File I want information from but don't want to include in the output
    - [ ] ^ This came up as a part of the SBOM cataloger
    - [ ] Is this a scope image where they want their changes and not the base image?
      - [ ] Diff the two images and maybe do the scan on that?
    - [ ] Current implementation is all layers vs squashed. Is there an intermediate where we only use USER defined scope.
    - [ ] This ties into multiple sources as saying source is these layers rather than the image string
- [ ] Docker daemon provider arch variant https://github.com/anchore/stereoscope/issues/143
  - [ ] M1 mac user to validate when we fix the empty string for variant in stereoscope

---

15 September 2022

- [ ] Linked SBOM discovery via rekor https://github.com/anchore/syft/issues/1159
    - [ ] We'll chat about this in the next meeting
- [ ] Better detection of dependencies in Official Docker Images https://github.com/anchore/syft/issues/1197
    - [ ] Hash lookup, look at known base images for known versions. There are two flavors of this, one on rekor the other using a internal set
    - [ ] Run --version on the binary in a container
    - [ ] Can we attach this kind of information into the image itself in the OCI registry?
        - [ ] Buildkit issue for reference types https://github.com/moby/buildkit/issues/2773
    - [ ] Could we start attaching some of this information for binaries as suffixes in the binary?
        - [ ] Like mac signing
        - [ ] ELF metadata: http://systemd.io/ELF_PACKAGE_METADATA/
- [ ] No packages discovered in SIF when image source not specified https://github.com/anchore/syft/issues/1189
    - [ ] https://github.com/anchore/stereoscope/pull/142 <- can't seem to join Zoom call due to outage, but looking for guidance/code review, thanks (Ad Hughes)
- [ ] I just filed this issue, I'd appreciate feedback: https://github.com/anchore/syft/issues/1207

---

1 September 2022

Topics (*no specific order, anyone can add topics! Items marked off means they have been covered in conversation*):

- [x] ~~PR for SBOM cataloger addition https://github.com/anchore/syft/pull/1029~~
    - [ ] Source and distro updates needed before we can proceed with the SBOM cataloger
    - [ ] https://github.com/anchore/grype/issues/562 and https://github.com/anchore/syft/issues/562 are just a couple related to this
- [x] ~~Bring in GitLab Community Advisories as an additional data source and surface a flag to disable CPE-based matching by default within Grype~~
    - [ ] Adding the gitlab driver is anchore internal processes - some OSS sync on the grype consumption side but majority of work will be on the internal builder side
- [x] ~~Almost duplicate packages being found by catalogers with internal overlap~~

        - [ ] https://github.com/anchore/syft/issues/1162
        - [ ] Remove location field from ID Hash generation
        - [ ] Merge location data between identical packages
        - [ ] Virtual Path field in Metadata? Do we remove this? Or turn it into a list and merge it downstream
- [x] ~~Grype database missing CVSS vectors for GHSA https://github.com/anchore/grype/issues/904~~
    - [ ] Update Feed Driver
    - [ ] Then update the database builder to pull in driver changes
- [x] ~~Enable formatters to take in external library structs https://github.com/anchore/syft/pull/1172~~
- [x] ~~PR to fix SIF packages when image source not specified https://github.com/anchore/syft/issues/1189~~, https://github.com/anchore/stereoscope/pull/142

Notes:

- PR for SBOM cataloger
    - Need to support multiple sources and distros in the syft json output
        - Add distro as an artifact?
    - Allow for SBOM merge strategies
- GitLab data source
    - Add configurable in grype for which namespaces to allow/deny
- Duplicate packages
    -

---

18 August 2022

Topics (*no specific order, anyone can add topics! Items marked off means they have been covered in conversation*):

- [ ] Merge duplicate packages https://github.com/anchore/syft/issues/1162
- [ ] Online verification of artifacts: https://github.com/anchore/syft/issues/1115

---

4 August 2022

Topics (*no specific order, anyone can add topics! Items marked off means they have been covered in conversation*):

- ☑ ~~General question about missing versions from purls/sbom packages from @deseena~~
- ☑ ~~https://github.com/anchore/grype/issues/837~~
    - Can the docker image be republished whenever a new Grype db is created in s3?
    - Use **grype:ci** tag where the container holds the most recent version of Grype and the newly created database?
    - Example workflow https://github.com/anchore/grype/pull/841
- ☑ ~~NPM Licenses~~
    - How to solve / philosophies of syft
    - Summary
        - When looking at additional locations, any file that is opened and somehow makes it into the SBOM should be listed in the “locations” field on the package.
- ☐ Online verification of artifacts: https://github.com/anchore/syft/issues/1115

---

21 July 2022

Topics (*no specific order, anyone can add topics! Items marked off means they have been covered in conversation*):

- ☑ ~~https://github.com/anchore/syft/issues/845~~
    - License population for directory scans
- ☑ ~~https://github.com/anchore/syft/issues/1113~~
    - License aware analysis from manifest
    - Easy since node_modules are sibling folder
    - Should we just do it from package.json?
    - Do we have to do verification on modified modules vs manifest files?
    - Start simple and get catalogers to search for other sibling manifest files for more metadata to merge information
    - Summary - we are ok with expanding syft to include more ecosystem aware cataloging for single file inputs
- ☑ ~~https://github.com/anchore/grype/pull/795~~
    - New line separate purl list
    - Rename purl input
- ☑ ~~https://github.com/anchore/grype/issues/800~~
    - Grype is shadowing the Redis databases's vulnerabilities over the pypi redis package
    - Pinged weston and covered Pull Request to explore doing these corrections at the data layer
- ☑ ~~https://github.com/anchore/syft/issues/1066~~

    - Syft generates cpes for python redis that shadow the real redis cpes
    - See above comment about incorporating these corrections at the data layer
- ☑ ~~https://github.com/anchore/syft/issues/933~~
    - License processing when converting to report format
- ☐ https://github.com/anchore/syft/issues/660
    - License swap to no assertion
- ☑ ~~syft command to show globs and matching pattern for different catalogers~~
    - Should we allow users to tinker with the lists of matches for specific catalogers?
    - Syft catalogers <directive> to show current catalogers
    - If we lift this into the configuration then -vv just prints the default for people to observe

Action items:

- Add the above command to new issue for more transparency on glob matchers

---

7 July 2022

Attendees (*please add your name when you join in the zoom chat*):

- Christopher Phillips

Topics (*no specific order, anyone can add topics! Items marked off means they have been covered in conversation*):

- ☑ ~~https://github.com/anchore/syft/issues/833~~
    - Distroless static? I've only built from scratch, but the notes in the issue about gitlab CI and need for a shell on k8s based runners is interesting
        - Outcome commented on the issue (add a new image build)
- ☑ ~~Enhance behavior on syft scope:~~
    - https://github.com/anchore/syft/issues/15 - new scopes
    - https://github.com/anchore/syft/issues/435 - maybe configurations for scopes
    - Revisit package deduplication:
        - Should we remove location from the definition of a package ID? And if so, do we want to keep the concept of number of instances of a package? Or allow this to be lossy?
- ☐ pURL language extension
    - Should we make package language to a slice of languages?
    - If there are multiple languages then we can provide all possible languages in the case where we don't know the exact language
    - Maybe we should revisit how ecosystem vs language matching in grype works…
- ☐ Grype DB Diff PR Review/Overview

    - Add default behavior to the command? Diffing the last two that have been released.
    - There are some differences between the two latest DBs (in the hundreds)... this seems suspicous

Notes:

- Syft scope
    - Augment the scope selections to surface the file for a path that:
        - Is the first instance
        - Is continuous with the target layer (e.g. for squashed there is no removed file for this location between the top “left” file and the lowest layer found)
        - For selection of the earliest location for a package file, the package that is represented MUST have the same ID (content change detection)
    - We can change the default behavior here and optionally surface this as a configurable option
    - Change package ID to be entirely decoupled from location (finding two packages in different locations that are otherwise the same should be considered identical, thus have the same package ID)

Action items:

- ☑ ~~Talk with Weston and Josh about grype updates for language/ecosytem aware feeds.~~
    - Example: we run into JVM related issues where it could be multiple languages and syft only specifies one, how can grype recognize this and search all related feeds.
        - PURL related issue https://github.com/package-url/purl-spec/pull/178
        - The above issue links syft/grype related developments on package catalogers where multiple languages are possible

---

23 June 2022

Attendees (*please add your name when you join in the zoom chat*):

-

Topics (*no specific order, anyone can add topics! Items marked off means they have been covered in conversation*):

- ☐ https://github.com/anchore/syft/issues/833

    - Distroless static? I've only built from scratch, but the notes in the issue about gitlab CI and need for a shell on k8s based runners is interesting
- [x] ~~https://github.com/anchore/grype/issues/764~~
    - Follow up on grype DB Diff from last meeting
- [x] ~~https://github.com/anchore/grype/issues/679~~
    - Getting syft/grype into Gallery ECR AWS
- [x] ~~Scan multiple sources (may produce one or more SBOMs)~~ ~~https://github.com/anchore/syft/issues/562~~
- [x] ~~ECR Support for Syft/Grype~~
    - https://github.com/anchore/grype/issues/774
    - https://github.com/anchore/stereoscope/issues/65

Notes:

-

Action items:

- [ ] Comments for https://github.com/anchore/grype/pull/789/files:
    - Drop generics
    - Keep old linter version
    - Try not to use the serialized comparison (no need to change the interface)
    - Alex Goodman provide detailed comments
- [ ] Comments for https://github.com/anchore/stereoscope/pull/131:
    - Foreshadow: GCR auth might change
    - Azure doesn't need to be supported
    - Alex Goodman provide detailed comments
- [ ] Multiple sources:
    - Take the relationships approach
    - Maybe make source-to-X relationships optional (via configuration)
    - How will this work with spdx or cyclonedx? (they don't allow for multiple source blocks / descriptor blocks)
        - SPDX: maybe in the next release? For now we only have the document namespace
            - [ ] For the meantime, maybe we error out when attempting to generate a merged document that is an SPDX format?
        - CycloneDX: we can get away with multiple component entries
        - Github format: probably not supported
            - [ ] Also a case to error out
    - Consider not even allowing for multiple inputs for a single syft call, instead add a merge command
        - Could we provide a "split" command to separate back out a merged sbom?

---

9 June 2022

Attendees (*please add your name when you join in the zoom chat*):

-

Topics (*no specific order, anyone can add topics! Items marked off means they have been covered in conversation*):

- ☑ ~~https://github.com/anchore/grype/issues/739~~
    - Looking to add support to show layers where vulnerabilities are located
- ☐ https://github.com/anchore/syft/issues/833
    - Distroless static? I've only built from scratch, but the notes in the issue about gitlab CI and need for a shell on k8s based runners is interesting
- ☐ https://github.com/anchore/grype/issues/764
    - Follow up on grype DB Diff from last meeting
- ☐ https://github.com/anchore/grype/issues/679
    - Getting syft/grype into Gallery ECR AWS
- ☐ Scan multiple sources (may produce one or more SBOMs) https://github.com/anchore/syft/issues/562
- ☐ ECR Support for Syft/Grype
    - https://github.com/anchore/grype/issues/774
    - https://github.com/anchore/stereoscope/issues/65
- ☑ ~~SBOM cataloger~~
    - https://github.com/anchore/syft/pull/1029
        - Note: right now this approach will merge SBOM findings in without noting where they were they came from
        - Only packages are kept, however, all other artifacts (linux distro, relationships, etc) are lost.
    - https://github.com/anchore/syft/issues/737
    - https://github.com/anchore/syft/pull/1029#issuecomment-1146872960
      "A larger question re:design is that we have now introduced a dependency on our internal formatting library on the cataloger. The way that syft is currently structured, we have kept the formatters as a separate layer of logic that build on top of the syft sbom model. This is probably something we might want to think about as we introduce dependencies between the formatters and the catalogers.

> The PR also currently only parses cyclonedx SBOMs, we should probably extend it to the other formats syft supports if we agree with the proposed design.
>
> One of the other sticking points we have discussed in the past when we have talked about cataloging sboms is providing links back to the original sbom as evidence. Currently syft is able to provide details as to 'why' it included a specific component in the output. One of the blockers to implementing #737 has been figuring out an appropriate way to represent that syft itself didn't discover the cataloged components, rather it pulled them from an sbom present in the image”

Notes:

- SBOM Cataloger:
  - Consider extending the syft model with SBOM components
  - Should we omit scanning certain globs / locations based on information found within other SBOMs?
    - https://github.com/anchore/syft/issues/31
    - Maybe in another file
    - Issue 31 and 737 should be considered together in terms of use cases --this will help get them unstuck

Action items:

- [ ] Capture thoughts for format and cataloger injection via application config in https://github.com/anchore/syft/issues/558 Alex Goodman (see https://gist.github.com/samj1912/91be21fb427a42a1f98197667a3b047e)
- [ ] Extension point: add formats as injection points
- [ ] Extension point: add package catalogers
- [ ] Extension point: filter sbom.Artifacts (maybe expose tasks)
- [ ] Alex Goodman Christopher Phillips update syft 1029 with context

---

24 May 2022

Attendees (*please add your name when you join in the zoom chat*):

-

Topics (*no specific order, anyone can add topics! Items marked off means they have been covered in conversation*):

- [ ] Enable/Disable catalogers via CLI and config:
    - https://github.com/anchore/syft/pull/888
    - https://github.com/anchore/syft/pull/843
    - https://github.com/anchore/syft/issues/840
- [x] ~~Singulatiry Image Support~~
    - https://github.com/anchore/stereoscope/pull/125
    - See PR for progress on hack out GPL library
- [x] ~~Turn on other non-package catalogers by default~~
    - https://github.com/anchore/syft/issues/477 (file metadata)
    - Consider enabling the file digests cataloger by default
    - Collect all file metadata for all files on system (already implemented, off by default)
    - Can lead to large sboms
    - Windows is showing huge SBOM so worth keeping our finger in the air so we maintain parity of this becomes a standard
- [ ] Scan multiple sources (may produce one or more SBOMs) https://github.com/anchore/syft/issues/562
- [x] ~~We should start using the upstream go-rpmdb fork to take advantage of the sqlite3 support added. There is also community pressure to do so - https://github.com/knqyf263/go-rpmdb/pull/19~~
- [x] ~~SBOM Centric Workflow for License -~~
    - https://github.com/anchore/syft/issues/933
    - https://github.com/anchore/syft/issues/656
    - Follow up on both specifications for license declarations
    - https://github.com/google/go-licenses
    -
    - Plans for CVE feeds in the future [Chapman Pendery]

Notes:

- Hack out the GPL license library
    - `Unallowable license (GPL-2.0) from "github.com/rasky/go-lzo"`
    - `Unallowable license (CC0-1.0) from "github.com/therootcompany/xz`
- Grype Diff command for Database
    - No current way to have multiple db installed
    - Grype diff would show what is new (any adds) and what's updated (any updates)
    - Db diff for two points in time
    - Usecase see what's new and see what's changed
        - Output would need to be machine readable (also possible default table output)

    - `grype db` current commands
        - check check to see if there is a database update available
        - delete delete the vulnerability database
        - import import a vulnerability database archive
        - list list all DBs available according to the listing URL
        - status display database status
        - update download the latest vulnerability database
- License:
    - Where we don't get provided licenses we can still start reading them to detect and surface them to the top of the SBOM
    - SBOM output view for licenses:
        - For every package metadata do we have license metadata? (license as part of package cataloger)
        - OR - does it make sense as a file cataloger
            - Output shows THESE FILES have licenses output and we form a relationship with package

Action items:

- [ ] Add cco-1.0 to allow list on bouncer.
- [ ] Clean up / investigate detection issues around https://github.com/anchore/syft/issues/933 so all formats are reporting licenses accurately
- [ ] Remove RPMDB fork and move upstream
- [ ] Find time to get Enable Disable catalogers across finish line for PR and issue differences

---

12 May 2022

****Canceled****

---

28 April 2022

Attendees (*please add your name when you join in the zoom chat*):

- Alex Goodman
- Adam Hughes (Sylabs)
- Zach Hill
- Keith Cunningham
- Jon Velando (rigzba21)
- Sambhav Kothari (Bloomberg)
- Christopher Phillips
- Keith Zantow
- Jonas Galvao Xavier
- Weston Steimel

Topics (*no specific order, anyone can add topics! Items marked off means they have been covered in conversation*):

- [ ]

Notes:

-

Action items:

- [x] ~~VEX + Syft compatibility~~
    - https://github.com/anchore/grype/pull/678
    - SBOM Ref / link:
        - Current SBOM model in syft doesn't have the original metadata from CycloneDX (such as UUID for the BOM ID) – can we persist this for later transformations?
            - [ ] We could keep the format objects around on sbom.SBOM or in another object for every package
            - [ ]
        - Syft decode/encode cycles not reserving the package ID
            - [ ] A PR has fixed this a week ago
        - Could this be informed by the syft issue for capturing/cataloging info from discovered SBOMs?
            - [ ] There could be overlapping mechanisms there
        - Usecase: grype ingests a list of SBOMs… how does the reporting work?

            - [ ] Does the list come from a single source? Or provided individually (as siblings)?
        - Grype should probably fail when the input format specified cannot generate a vex embedded output. We should not default to always generating a vex document.
        - If Grype is given the lose umbrella of wanting a Vex document it can be smart enough to discover if it should be embedded or not for the user (depends on a umbrella “vex-xml” and “vex-json” formats which are “smart”... other more explicit ones error out)
- [ ] Python conda (updating egg/wheel cataloger)
- [ ] Enable/Disable catalogers via CLI and config:
    - https://github.com/anchore/syft/pull/888
    - https://github.com/anchore/syft/pull/843
    - https://github.com/anchore/syft/issues/840
    - Open questions:
        - Current suggestions are for package catalogers only, however should we also include non-package catalogers?
        - How should groups of catalogers work? (config profiles, lists, etc)
    - Notes:
- [ ] Turn on other non-package catalogers by default
    - https://github.com/anchore/syft/issues/477 (file metadata)
    - Consider enabling the file digests cataloger by default
- [ ] Scan multiple sources (may produce one or more SBOMs) https://github.com/anchore/syft/issues/562
- [x] ~~Capture shared libraries~~ ~~https://github.com/anchore/syft/issues/661~~
    - We should consider when to catalog these based on the source being scanned (maybe images and dir only? Maybe not individual files?)
- [x] ~~Singularity image support in Stereoscope/Syft~~
    - https://github.com/anchore/syft/issues/937
    - https://github.com/anchore/stereoscope/pull/125
    - Notes:
        - There are some readers not closed Alex Goodman take a closer look at open Fds / image cleanup path

---

14 April 2022

Attendees (*please add your name when you join in the zoom chat*):

- Christopher Phillips

- Jon Velando

Topics (*no specific order, anyone can add topics! Items marked off means they have been covered in conversation*):

- ☑ ~~Conda support~~ ~~https://github.com/anchore/syft/issues/932~~
  - ☐ https://docs.conda.io/en/latest/

    Open Questions:
    - Can it fit under the current python cataloger or should it be separate? A meta cataloger that composes features from language specific catalogers. (Python, R, Ruby, Lua, Scala, Java, JavaScript, C/ C++, FORTRAN)
    - Conda separate metadata (is this per language cataloger)?
    - Are direct and transitive already identified through current pip identification or does conda do this another (better?) way?
    - Identification of channels?
      - https://docs.conda.io/projects/conda-build/en/latest/concepts/generating-index.html
    - Upstream checks or all on disk identification?
    - Inclusion of tarball specific data (can this vary from machine to machine). If someone has brought their own package with updated non cannon metadata (numpy ⇒ numnum), but the digest for that code matches some known package can/should we link those in some way?
    - How does grype do the vulnerability matching against a “conda” specific
    - Ways to install conda itself. Self extracting bundle bash script to get all base conda environment.
      - https://github.com/conda-forge/miniforge

Notes:

- Current package managers pick up from wrong source (think it’s pip or NPM but is conda). Missing dynamic libraries or shared object files. C and C++ packages are totally missing.
- Should pip or NPM catalogers be made smarter in the ability to discover sub language dynamic libraries as seen above
- Ops has to provide manifests to the data scientists and a distinction of these packages
- Log stash Example - contained all the languages - When dealing with updates packages with same name (ruby and JAVA). Syft can tell which language
- Lead with language and then tell how we discovered it
- Currently assigned and discovered under egg_wheel cataloger
- Active base environment has certain env variables (LOOK THESE UP) syft can key on these to help establish the shape of the conda packages without doing any kind of network request or base environment configuration
  - CONDA_EXE - path to the conda executable
  - CONDA_PREFIX - path to the (base) environment
  - CONDA_DEFAULT_ENV - name of the default conda environment

- https://github.com/anchore/syft/blob/main/syft/pkg/cataloger/python/package_cataloger.go Current detection when conda exists but picked up by egg_wheel cataloger
- Follow up on current issue after talking with alex based on this feed on direction we would like to go. Outline architecture as far as package discovery and conda metadata integration in python package cataloger.

- [x] ~~Enable/Disable catalogers via CLI and config:~~
    - [ ] https://github.com/anchore/syft/pull/888
    - [ ] https://github.com/anchore/syft/pull/843
    - [ ] https://github.com/anchore/syft/issues/840
    - [ ] Open questions:
        - [ ] Current suggestions are for package catalogers only, however should we also include non-package catalogers?
        - [ ] How should groups of catalogers work? (config profiles, lists, etc)
    - [ ] Notes:
- [ ] Turn on other non-package catalogers by default
    - [ ] https://github.com/anchore/syft/issues/477 (file metadata)
    - [ ] Consider enabling the file digests cataloger by default
- [ ] VEX + Syft compatibility
    - [ ] Notes:
        - [ ] Follow up with Sambhav in slack since he had to drop
- [ ] Community feedback on Docker SBOM
    - [ ] https://www.docker.com/blog/announcing-docker-sbom-a-step-towards-more-visibility-into-docker-images/
- [x] ~~SIF Issue ->~~
    - [ ] https://github.com/anchore/syft/issues/937
    - [ ] https://github.com/anchore/stereoscope/pull/123
    - [ ] Current PR does mount - but read squash FS is on a fork that implements OCI layer
    - [ ] Notes:
        - [ ] Read squash FS files
        - [ ] Mount is needed and shell out to mount needed for one of the solutions

Notes:

- JOSH: Seeing grumblings of converting from syft JSON → Something else
    - Syft convert syft.json → spdx/cyclone
    - EO does not specify format BTW
    - FDA medical device SBOM (SPDX)
    - Generate syft first allows us to output newer versions as we convert forward

Action items:

- Conda Support follow up using notes from above on issue with path forward for PYTHON specific implementation/direction for package cataloger and wheel discovery
- SYF follow up on pull request for library of non mount implementation
- Follow up on syft conversions from dense syft format → future proof of other sbom conversions (https://github.com/anchore/syft/issues/563

31 March 2022

Attendees (*please add your name when you join in the zoom chat*):

- Alex Goodman
- Jonas Galvao Xavier
- Frankie GJ
- Ryan Moran
- Sophie Wigmore
- Christopher Phillips
- Keith Zantow

Topics (*no specific order, anyone can add topics! Items marked off means they have been covered in conversation*):

- ☑ ~~Provider encoders/decoders for all major syft versions~~
    - https://github.com/anchore/syft/issues/846
    - Open Questions:
        - Should SPDX and CycloneDX have the same support assumptions for encoding and decoding as the Syft json format?
        - Should we export the format code as part of the public API?
            - Yes!
        - Proposal: Encoders for all major versions - Decoders for Latest of the Major versions
        - CycloneDX 1.3 support (difficult since cyclonedx consumer library does not support multiple minors)
    - Context Notes:
      Paketo Build packs and SBOM output context

        - Build packs uses syft as a Cli and Library
        - Creates a container image from source code
        - Here is source code ——> executable format (buildpack - user does not bring their own stuff)
        - SBOM we output as part of the API. Consumers can crab a container image and pull the SBOM off of it
        - Declare in the image labels syft version so consumers know how to decode the format. (STABILITY) PIN TO A CERTAIN SYFT VERSION
        - Syft that does scanning vs formatted output seam has been helpful
        - Missing the ability to specify different versions (syft 1 - 3)
        - Consumers use grype (needs specific syft SBOM format)

- Shipped version currently needs ability to specify syft version

☑ ~~Convert between multiple SBOM formats: https://github.com/anchore/syft/issues/563~~
  - Odds are that most conversions will be lossy if not starting with the syft-json format
  - Are there good ways to let users know of potential conversion problems? Or should we actively stop conversions that could have problems (allow for an override like `--force`)?
    - We probably do not need a detection mechanism, just a format-to-format pairing that would output a canned warning is good enough. A detection mechanism could be very difficult and it's not clear that it's a desired feature.
  - Notes:
    - Follow up with Hector on the issue since he does have a certain use case for conversion
    - When there is more distance between producers and consumers SBOMs become less valuable. Cost of regenerating vs the disruption of time between when a producer creates something and a consumer ingests it
    - Very useful for cases where you are building a polyglot app and need to merge dependencies from different stages of the pipeline (go vs javascript), but which end up in the same final application
    - Question: sbom merging capabilities
      - Would this work for multiple sources? Or only work for a single source?
      - This would at least be a syft schema change, we should explore what we would need for SPDX and CycloneDX

☐

Notes:

-

Action items:

☐

---

17 March 2022

Attendees (*please add your name when you join in the zoom chat*):

- Alex Goodman
- Weston Steimel
- Dan Luhring
- Josh Bressers
- Christopher Phillips
- Keith Zantow
- Jeff Buddington (JCI)
- Josh Knarr

Topics (*no specific order, anyone can add topics! Items marked off means they have been covered in conversation*):

- ☑ ~~Syft Keyless Attestation preview~~
  - Look more into what sigstore does for validations outside of our tooling
  - What policies exist that we can leverage so our tools and accept those to allow for vulnerability scanning on “verified” inputs
  - Latency and compute power for scanning is in high demand and a friction
  - Admission controller times out after 10 seconds
  - Output for the grype keyless attestation will be key for policy readers
- ☑ ~~Improve license detection in Syft:~~
  - Classify licenses from file content: https://github.com/anchore/syft/issues/656
  - Find SPDX License identifiers: https://github.com/anchore/syft/issues/565
  - We can start capturing full license text for files/packages and reporting in SPDX output (does CycloneDX have support for expressing this in output?)
  - This appears to be most useful for tracking IP via SBOMs, however, are the other use cases not captured here for this work?
- ☐ Convert between multiple SBOM formats: https://github.com/anchore/syft/issues/563
  - Odds are that most conversions will be lossy if not starting with the syft-json format
  - Are there good ways to let users know of potential conversion problems? Or should we actively stop conversions that could have problems (allow for an override like `--force`)?

Notes:

- Syft Licenses

    - Knowing that the license is different from a known license content is useful
    - Capturing license text is heavy, but possibly important, this could be opt in (helpful for follow up action for auditors)
    - Maybe include threshold information (if classifying) or a diff to known output if declared.
    - External license fetching, this could be interesting for change detection (if not just enrichment)
    -

Action items:

- [ ]

---

3 March 2022

Attendees (*please add your name when you join in the zoom chat*):

- Alex Goodman
- Keith Zantow
- Jeff Buddington (JCI)
- Christopher Phillips
- Jonas Galvao Xavier
- Dan Luhring
- Weston Steimel
- Alex Vanderpot

Topics (*no specific order, anyone can add topics! Items marked off means they have been covered in conversation*):

- [ ] Syft lib improvements
    - https://github.com/anchore/syft/issues/558
    - https://gist.github.com/wagoodman/57ed59a6d57600c23913071b8470175b

- [ ] Syft selectable / configurable catalogers
    - https://github.com/anchore/syft/issues/840
    - https://github.com/anchore/syft/pull/843
    - https://github.com/anchore/syft/issues/465

Notes:

- Syft lib improvements:
    - The source.Input is the right direction –having a string as input doesn't help out as much from an API point of view
    - Example: should we accept io.Reader instead of input paths from the filesystems
    - Consider funcional options for syft lib for cataloging
    - Possibly export CLI, app config, and other internals
        - Could expose via main() parameters / config
        - Expose cmd objects from cmd package (and move things out of internal)
        - Still don't export behavior (such as tasks) from the cmd package (and similar)
        - There are some parallels here with Format objects (accepting configuration)
    - Usecase: custom catalogers that append on to the typical execution path

- Side note: RFCs should go where? Issues? Hackmd.io?
    - See https://github.com/pivotal/kpack
    - https://github.com/pivotal/kpack
    - https://github.com/pivotal/kpack
- Syft selectable / configurable catalogers
    - Organize catalogers in a hierarchy (maybe not? See next point…)
    - Tags could be used on each cataloger to organize behavior without stovepiping to a single hierarchy
    - Consider using Tags() in the package cataloger interface

Action items:

---

17 Feb 2022

Attendees (*please add your name when you join in the zoom chat*):

- Alex Goodman
- Josh Bressers
- Keith Zantow
- Jeff Buddington (JCI)
- Sambhav Kothari (Bloomberg)
- Christopher Phillips
- Jonas Galvao Xavier

- Dan Luhring

Topics (*no specific order, anyone can add topics! Items marked off means they have been covered in conversation*):

- ☑ ~~Inclusion of CycloneDX 1.4 and VEX into syft / grype~~
    - https://github.com/anchore/syft/issues/744
    - https://github.com/anchore/grype/issues/591

Notes:

- CycloneDX 1.4 handy links:
    - https://cyclonedx.org/specification/overview/
    - https://cyclonedx.org/capabilities/
    - https://cyclonedx.org/docs/1.4/json/
    - https://github.com/CycloneDX/bom-examples
-

Action items:

- ☐

---

3 Feb 2022

Attendees (*please add your name when you join in the zoom chat*):

- Alex Goodman
- Keith Zantow
- Dan Luhring
- Weston Steimel
- Sambhav Kothari (Bloomberg)
- Josh Bressers

Topics (*no specific order, anyone can add topics! Items marked off means they have been covered in conversation*):

- ☑ ~~Syft - Add support for package dependency relationships~~

Notes:

- Reach out to CycloneDX folks about the future of relationships and edge qualifiers
- How might relationship types or degree differentiation help us with features in grype

- A minimum bar for us decoding wise: starting with (a syft generated doc) the target format, decoding to syft model, and encoding to back to the target format should result in the same document. We have an integration test in syft that does part of this –we should consider expanding on this.
- Should we look at Nix to get a “complex” example of how dependency descriptions would work?
- Looks like we’re generally in agreement with continuing with the package dependency work, but we should be careful with the sizing (needs to be split up) and balancing this work with features that improve the overall workflow (such as attestations).
- CycloneDX is thinking about adding an “evidence” field similar to our “Locations” field. We should vote on this! Checkout the specifications channel for the CycloneDX slack. https://github.com/CycloneDX/specification/issues/129
- Grype: consider adding a CVSS vector filter

Action items:

- [ ] Follow up with Nix cataloger to help out
- [ ] Next time we can meet and chat about CycloneDX 1.4 / VEX… raw thoughts / how that might fit into syft and grype
- [ ] Sam: New issue for CVSS vector filter

---

20 Jan 2022

Attendees (*please add your name when you join in the zoom chat*):

- Alex Goodman
- Christopher Phillips (Anchore)
- Jeff Buddington (JCI)
- Josh Bressers
- Sambhav Kothari (Bloomberg)
- Weston Steimel

Topics (*no specific order, anyone can add topics! Items marked off means they have been covered in conversation*):

- [ ] Grype - Guarantee more deterministic output
- [ ] Syft - Add support for package dependency relationships
- [x] ~~Syft - Sign and attest SBOM~~ + ~~https://github.com/anchore/syft/pull/759~~

Notes:

- Syft attestations possible command formats
    - syft attest -key $MY_PRIVATE_KEY dir:./ > attestation.json
    - syft attest –attach -key ubuntu:latest
    - possible future syft reading attestations from images specific to itself
- Open questions:
    - Should we only support images for now? And support other sources later? (this would be easier)
- “Keyless” signing via github support in v1 or fast follow to BYOK (Bring your own key)
- What other PKI infrastructure should be supported day 1
- If syft is to generate an sbom with a given format (SPDX or CycloneDX) then defer to those predicate types in the attestation rather than making anchor.spdx.v1
- Current examples of attestation validation via kyverno
    - https://kyverno.io/docs/writing-policies/verify-images/#verifying-image-attestations
    - https://github.com/JimBugwadia/image-verification-policy/blob/main/check-sbom.yaml
    - This is important because it’s something we want to be sure our attestations can plug into on day 1
- Topic for next time: support for package dependencies
    - Summary: today we raise up a lot of package manager information within the syftjson format, but is not used in other formats. We have the ability to start creating package-to-package relationships that describe structurally the dependency tree of packages that we discover. We want to explore this topic with the community (is this something you might find useful? What could this look like? etc)

Action items

- ☐ Christopher Phillips tag participants in v1 of the syft attestation PR

6 Jan 2022

Attendees (*please add your name when you join in the zoom chat*):

- Alex Goodman (Anchore)
- Sambhav Kothari (Bloomberg)
- Keith Zantow (Anchore)
- Christopher Phillips (Anchore)
- Jeff Buddington (JCI)

Topics (*no specific order, anyone can add topics! Items marked off means they have been covered in conversation*):

- ☑ ~~Syft - cataloger plugins~~ — ~~https://github.com/hashicorp/go-plugin~~
- ☑ ~~Syft - OCI output format~~

Notes:

- Cataloger plugins:
    - We've looked at 3 approaches in the past:
        - Go stdlib plugin system
            - Performant, though has caveats regarding app-to-plugin go versions
        - Host IPC support
        - Hashicorp plugin system
            - Lots of boilerplate but flexible
        - Simple stdout approach (like kubectl plugins)
    - Two use cases for this:
        - Upstream PR is difficult
        - Proprietary package types you want to catalog (can't be a PR)
    - Extension point 1
        - Idea: interface for changing packages: []packages -> []packages function
        - We can do this for other kinds of artifacts (distro, relationships, etc.)
    - Extension point 2: What does this mean for the formats themselves. Do we allow for format encoder extensions?
        - What about cases where you might want to mutate or redact information when encoding (e.g. adding get params to a pURL, but leaving the pURL otherwise intact)
    - Security point of view: we don't want folks getting unexpected execution
        - This should be opt in
        - There should be an allow list
        - Should we have an audit trail of removed, added, modified components from a plugin workflow?
- Syft OCI registry:
    - We could leverage the multi-output -o work:
        - `-o json=registry://my/image:latest`
    - To dan's point, we could let cosign continue to handle this
    - Could we import this functionality from cosign?
    - We should confirm is disk needed for syft-cosign interchange? Or can we pipe it?
        - Or we could extend cosign
    - We can at least document this interaction between syft and cosign

Action items:

- [ ] Alex Goodman add more details to the issue on what's been tried, use cases, and options moving forward
- [ ] Sambhav Kothari update the syft OCI registry issue with some of the conversation points and options.

---

23 Dec 2021

Attendees (*please add your name when you join in the zoom chat*):

- Alex Goodman (Anchore)
- Sambhav Kothari (Bloomberg)
- Dan Luhring (Anchore)
- Jonas Galvao Xavier (Anchore)
- Keith Zantow (Anchore)
- Weston Steimel

Topics (*no specific order, anyone can add topics! Items marked off means they have been covered in conversation*):

- ☑ ~~Syft - Support multiple BOM formats~~
- ☐ Syft - Describe multiple SBOM scan targets
- ☐ Grype - Guarantee more deterministic output
- ☑ ~~Syft - Improve CycloneDX output~~
- ☑ ~~Grype - Buildpacks SBOM follow up~~
- ☑ ~~Syft - handling arbitrary distros~~ / ~~Syft - Remove strong distro type~~
- ☐ Syft - cataloger plugins - https://github.com/hashicorp/go-plugin
- ☐ Syft - OCI output format

Notes:

- SBOM from buildpacks follow up https://github.com/anchore/grype/issues/520:
  - Cataloging from discovered SBOMs from a source. First option: make a cataloger that simply decodes discovered SBOMs and raises up the `pkg.Package`s found (with relationships for all packages found back to the SBOM where it was found in the source). This would key off of the SBOM file extension. This would be a good separate issue to track 👍
  - From discussion: The suggestion of scheme addition `<ggcr-source>+<sbom-attachment-type>` seems like a good idea and maps well into the other features we're interested in (e.g. cosign)
  - We may be able to use GGCR (as buildpacks does) for SBOM blob fetching without fetching the entire image (a stereoscope change).
- Improve CycloneDX output
  - Idea: use specific struct tags for format-to-format mapping of fields
  - Revisit the idea of lossless format conversion:
    - For SPDX: https://github.com/anchore/syft/pull/578
    - We could do the same thing for CycloneDX with the “property” fields
- Distro type

    - Question: how should custom distros work? Possibly input from the user
    - This is most useful in crafting pURLs
    - Consider adding –distro-hint flag for a fallback
        - SYFT_DISTRO_HINT=mydebian

Action items:

- [x] ~~Alex Goodman Open up an issue to discuss options for lossless format conversion (https://github.com/anchore/syft/issues/723)~~
- [ ] Alex Goodman ~~Open an issue for a distro cataloger and tie together the existing issues (maybe take app config as options)~~ This doesn't make sense since package cataloging depends on distro results. This also doesn't get to the problem brought up by hinting with distro type.
- [x] ~~Alex Goodman Open up issue for distro hinting (https://github.com/anchore/syft/issues/736)~~

---

9 Dec 2021

Attendees (*please add your name when you join in the zoom chat*):

- Alex Goodman
- Chris Phillips
- Dan Luhring
- Keith Zantow
- Weston Steimel
- Sambhav Kothari
- Jonas Xavier
- Chet Burgess
- Josh Bressers
- Gabe Cemaj

Topics (*no specific order, anyone can add topics! Items marked off means they have been covered in conversation*):

- [x] ~~Grype - Add support for reading SBOM on image attached by buildpacks~~
- [x] ~~Grype - Reading image attached SBOM in cosign format~~ ~~<-- Related to buildpacks since they source from the same format~~
- [ ] Syft - Support multiple BOM formats
- [x] ~~Syft - OCI output format~~ ~~<-- Also related to 1 and 2~~
- [ ] Syft - Describe multiple SBOM scan targets
- [ ] Grype - Guarantee more deterministic output

Notes:

- Grype - Add support for reading SBOM on image attached by buildpacks
    - How do we find the layers related to the SBOM, what parsing is needed? Answer: we can glob based on the filename within the contained folder structure
    - Hint: provide syft or grype the dir directly for intaking multiple SBOMs (dir is from the SBOM layer)
    - Question: do we like the idea of using schemes for hinting at input and how to get/interpret information (e.g. SBOMs from buildpacks vs cosign attach)
        - We can use the digest directly without scheme and let what is there guide how to parse the information
        - Note for cosign: https://github.com/sigstore/cosign/blob/main/specs/SBOM_SPEC.md#scopes
    - Question: based on the last question, what is good default behavior?
        - Example: grype <repo>/<image>:
            - Look for a cosign attestation first
            - Look for a buildpack sbom second
            - Grab the full image last
        - Maybe exclude default options by configuration
    - Question: based on the last two questions, what is the best way to include “where did this come from” in the output in grype/syft?
        - A few things… one of these could be leveraged to help clarify output
            - there is a location field that we could leverage
            - We also have a source field at the root of the doc
        - The oddity about using source is that you may discover an SBOM within another source (say an SBOM within an image)... we should distinguish this on a per-package basis
        - Probably describe globally and link on a per-package basis
    - Observation: other than output we also have a need to have strict input validation to control risk (e.g. must come from a signed image, must come from an attestation, etc)
    - Workflow: having a document from vX of syft and you rescan with vY of syft, how do you reconcile the differences?
        - Question 1: how do you see what the differences even are?
        - Question 2: how do you know why the differences are there?

Action items:

- [ ]

<TEMPLATE>

Attendees (*please add your name when you join in the zoom chat*):

-

Topics (*no specific order, anyone can add topics! Items marked off means they have been covered in conversation*):

- [ ]

Notes:

-

Action items:

- [ ]