

Practical Software Supply Chain Security

Mike Vainio


Family stats:

2 kids
4 cats
1 wife




Current State of Software Supply Chain Security

Software Supply Chain



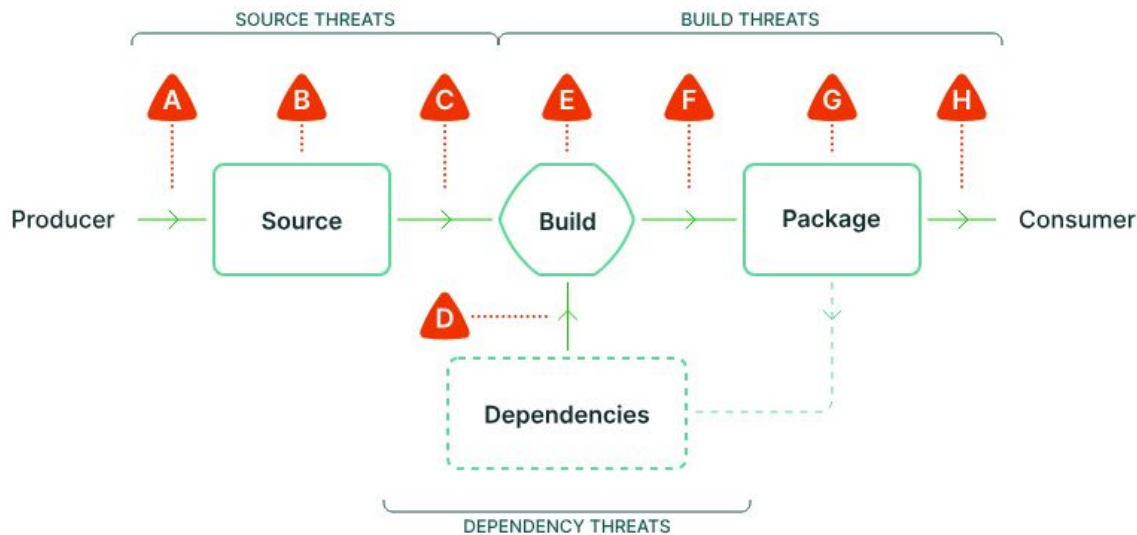
Supply chain attacks are becoming more common

FIGURE 1.7. NEXT GENERATION SOFTWARE SUPPLY CHAIN ATTACKS (2019-2023)



Source: <https://www.sonatype.com/state-of-the-software-supply-chain/open-source-supply-and-demand>

Threats all over the place



SOURCE THREATS

- A** Submit unauthorized change
- B** Compromise source repo
- C** Build from modified source

DEPENDENCY THREATS

- D** Use compromised dependency

BUILD THREATS

- E** Compromise build process
- F** Upload modified package
- G** Compromise package registry
- H** Use compromised package

Threats #2

	Integrity threat	Known example	How SLSA can help
A	Submit unauthorized change (to source repo)	SushiSwap : Contractor with repository access pushed a malicious commit redirecting cryptocurrency to himself.	Two-person review could have caught the unauthorized change.
B	Compromise source repo	PHP : Attacker compromised PHP's self-hosted git server and injected two malicious commits.	A better-protected source code platform would have been a much harder target for the attackers.
C	Build from modified source (not matching source repo)	Webmin : Attacker modified the build infrastructure to use source files not matching source control.	A SLSA-compliant build server would have produced provenance identifying the actual sources used, allowing consumers to detect such tampering.

D	Use compromised dependency (i.e. A-H, recursively)	event-stream : Attacker added an innocuous dependency and then later updated the dependency to add malicious behavior. The update did not match the code submitted to GitHub (i.e. attack F).	Applying SLSA recursively to all dependencies would have prevented this particular vector, because the provenance would have indicated that it either wasn't built from a proper builder or that the source did not come from GitHub.
E	Compromise build process	SolarWinds : Attacker compromised the build platform and installed an implant that injected malicious behavior during each build.	Higher SLSA levels require stronger security controls for the build platform , making it more difficult to compromise and gain persistence.
F	Upload modified package (not matching build process)	CodeCov : Attacker used leaked credentials to upload a malicious artifact to a GCS bucket, from which users download directly.	Provenance of the artifact in the GCS bucket would have shown that the artifact was not built in the expected manner from the expected source repo.

...

Software Bill of Materials

SBOM can help us understand what goes into an artifact

What's an SBOM?

```
[mike@linukka ~]$ syft nginx:alpine
```

✓ Loaded image

✓ Parsed image

✓ Cataloged packages

[66 packages]

NAME	VERSION	TYPE
alpine-baselayout	3.4.3-r1	apk
alpine-baselayout-data	3.4.3-r1	apk
alpine-keys	2.4-r1	apk
aom-libs	3.6.1-r0	apk
apk-tools	2.14.0-r2	apk
brotli-libs	1.0.9-r14	apk
busybox	1.36.1-r2	apk
busybox-binsh	1.36.1-r2	apk
ca-certificates	20230506-r0	apk

```
[mike@linukka verinotes]$ syft packages file:go.mod
```

✓ Indexed file system

✓ Cataloged packages

[29 packages]

NAME	VERSION	TYPE
ariga.io/atlas	v0.9.1-0.20230119145809-92243f7c55cb	go-module
entgo.io/ent	v0.11.8	go-module
github.com/agext/levenshtein	v1.2.1	go-module
github.com/apparentlymart/go-textseg/v13	v13.0.0	go-module
github.com/davecgh/go-spew	v1.1.1	go-module
github.com/go-chi/chi/v5	v5.0.8	go-module
github.com/go-chi/httplog	v0.2.5	go-module

```
1  {
2    "spdxVersion": "SPDX-2.3",
3    "dataLicense": "CC0-1.0",
4    "SPDXID": "SPDXRef-DOCUMENT",
5    "name": "go.mod",
6    "documentNamespace": "https://anchore.com/syft/file/go.mod-f518a5fa-36fc-4f90-996d-ec1a9e3b7c14",
7    "creationInfo": {
8      "licenseListVersion": "3.21",
9      "creators": [
10       "Organization: Anchore, Inc",
11       "Tool: syft-0.86.1"
12     ],
13     "created": "2023-11-20T12:39:12Z"
14   },
15   "packages": [
16     {
17       "name": "ariga.io/atlas",
18       "SPDXID": "SPDXRef-Package-go-module-ariga.io-atlas-4ba436b91a46da78",
19       "versionInfo": "v0.9.1-0.20230119145809-92243f7c55cb",
20       "downloadLocation": "NOASSERTION",
21       "filesAnalyzed": false,
22       "sourceInfo": "acquired package info from go module information: /go.mod",
23       "licenseConcluded": "NOASSERTION",
24       "licenseDeclared": "NOASSERTION",
25       "copyrightText": "NOASSERTION",
26       "externalRefs": [
27         {
28           "referenceCategory": "PACKAGE-MANAGER",
29           "referenceType": "purl",
30           "referenceLocator": "pkg:golang/ariga.io/atlas@v0.9.1-0.20230119145809-92243f7c55cb"
31         }
32       ]
33     }
34   ]
35 }
```

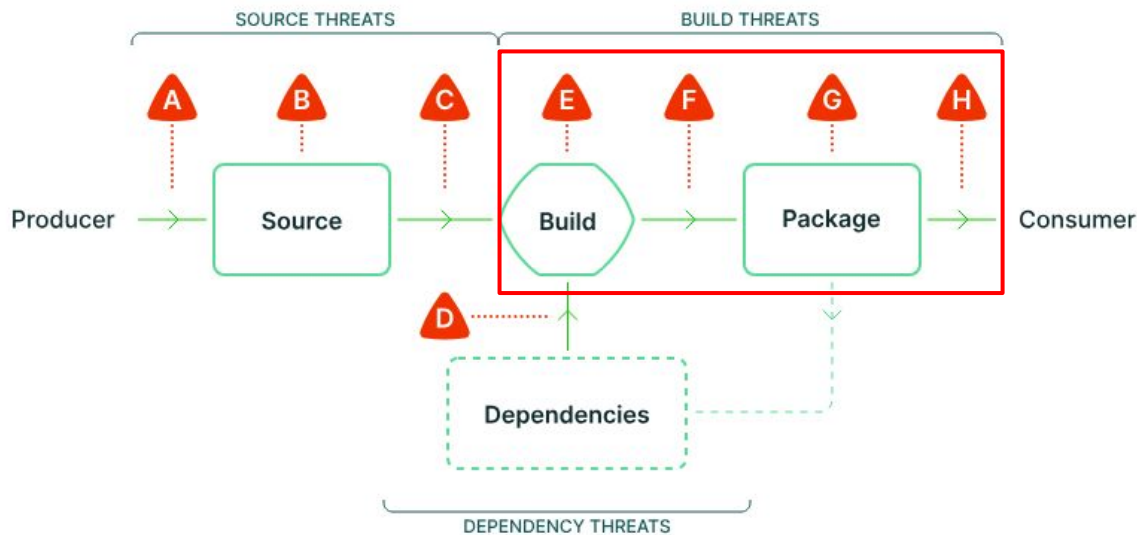
Supply-chain Levels for Software Artifacts

SLSA is a framework for software supply chain security

SLSA is not just about providing this metadata, it's also about treating your build system as a production system

SLSA helps to trace an artifact back to its source

SLSA v1.0 - Build Track



SLSA Security Levels (for Build track)

Track/Level	Requirements	Focus
Build L0	(none)	(n/a)
Build L1	Provenance showing how the package was built	Mistakes, documentation
Build L2	Signed provenance, generated by a hosted build platform	Tampering after the build
Build L3	Hardened build platform	Tampering during the build

SLSA Provenance

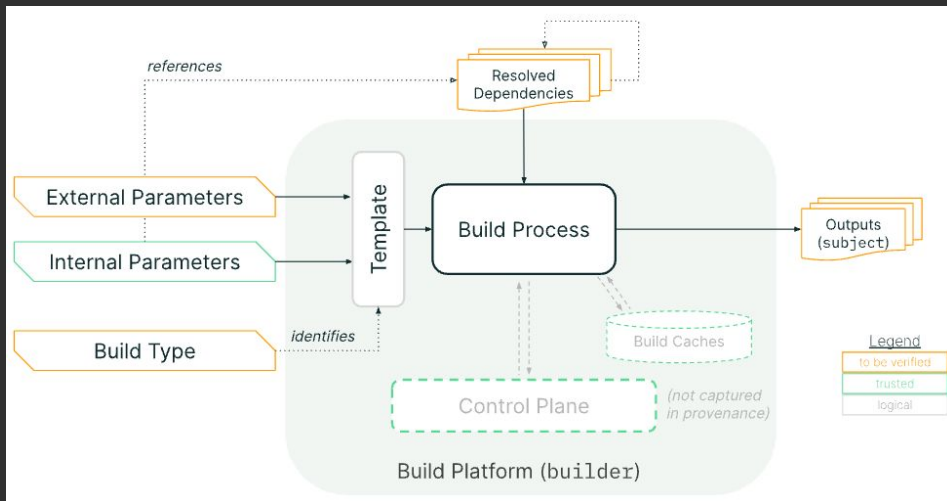
Provenance according to Google search:

the place of origin or earliest known history of something

SLSA provenance (according to SLSA.dev):

verifiable information about software artifacts describing where, when and how something was produced

SLSA Provenance



Source:

<https://slsa.dev/spec/v1.0/provenance>

```
1 {
2   "builder": {
3     "id": "https://github.com/slsa-framework/slsa-github-generator/.github/workflows/
4   },
5   "buildType": "https://github.com/slsa-framework/slsa-github-generator/container@v1"
6   "invocation": {
7     "configSource": {
8       "uri": "git+https://github.com/verifa/verinotes@refs/tags/v0.2.0",
9       "digest": {
10        "sha1": "d0a495e16d32c2bb125bf1481f747e711714199b"
11      },
12      "entryPoint": ".github/workflows/release.yaml"
13    },
14    "parameters": {},
15    "environment": {
16      "github_ref": "refs/tags/v0.2.0",
17      "github_ref_type": "tag",
18      "github_repository_id": "603037671",
19      "github_repository_owner": "verifa",
20      "github_sha1": "d0a495e16d32c2bb125bf1481f747e711714199b"
21    }
22  },
23  "materials": [
24    {
25      "uri": "git+https://github.com/verifa/verinotes@refs/tags/v0.2.0",
26      "digest": {
27        "sha1": "d0a495e16d32c2bb125bf1481f747e711714199b"
28      }
29    }
30  ]
31 }
```

How many here generate SBOMs?

How many here generate SLSA provenance?

75%

of leaders reported generating SBOMs for their applications

25%

of engineering professionals reported generating SBOMs for their applications



Sigstore

The Sigstore framework and tooling empowers software developers and consumers to securely sign and verify software artifacts

The project is backed by the Open Source Security Foundation (OpenSSF) under the Linux Foundation

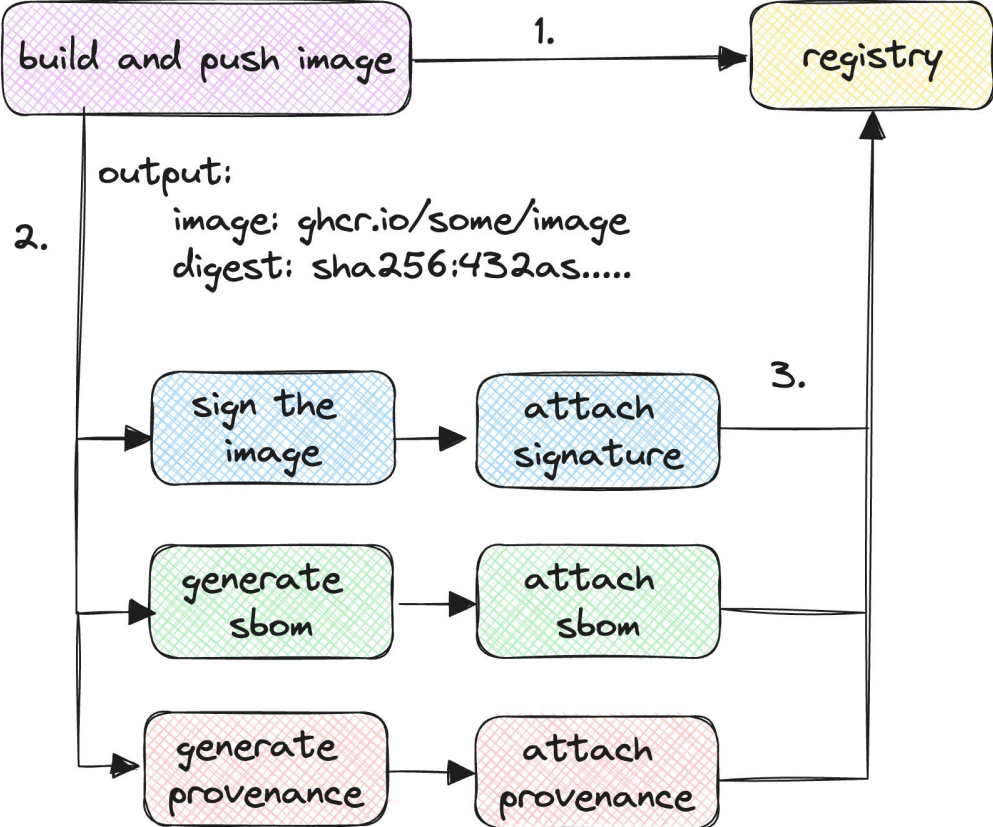


Main components:

- **cosign** CLI
- **Rekor** transparency log
- **Fulcio** code signing Certificate Authority.

Let's build!

Pipeline Overview



Inside the Registry

registry

tags for image:

→ v1.0.1

sha256:<hash>.sig

sha256:<hash>.att

Pipeline – Build and Push

```
2 name: Release workflow
3 on:
4   push:
5     tags:
6       - "*"
7
8 env:
9   KO_DOCKER_REPO: "ghcr.io/${{ github.repository }}"
10
11 jobs:
12   build:
13     runs-on: ubuntu-latest
14     permissions:
15       contents: read
16       packages: write
17     outputs:
18       image: "${{ steps.build.outputs.image }}"
19       digest: "${{ steps.build.outputs.digest }}"
20     steps:
21     - name: checkout repo
22       uses: actions/setup-go@fac708d6674e30b6ba41289acaab6d4b75aa0753 #v4.0.1
23     - uses: ko-build/setup-ko@ace48d793556083a76f1e3e6068850c1f4a369aa #v0.6
24
25     - name: Build and push with ko
26       id: build
27       run: |
28         # Build & push the image. Save the image name & digest
29         image_and_digest=$(ko build --tags="${{tag}}" --bare --sbom=none .)
30
31         # Output the image name and digest so we can generate provenance.
32         digest=$(echo "${image_and_digest}" | cut -d'@' -f2)
33
34         # digest/hash and image to outputs
35         echo "digest=$digest" >> "$GITHUB_OUTPUT"
36         echo "image=$KO_DOCKER_REPO" >> "$GITHUB_OUTPUT"
```

Pipeline - Sign Image

```
38   sign-image:
39     runs-on: ubuntu-latest
40     needs: [build]
41     permissions:
42       packages: write
43       id-token: write
44     env:
45       image: ${ needs.build.outputs.image }
46       digest: ${ needs.build.outputs.digest }
47     steps:
48     - name: Install cosign
49     - name: Login to ghcr.io
50
51     - name: Sign image
52       run: |
53         cosign sign "${image}@${digest}" --yes
```

Pipeline – Attach SBOM

```
55 sbom:
56   runs-on: ubuntu-latest
57   needs: [build]
58   permissions:
59     packages: write
60     id-token: write
61   env:
62     image: ${ needs.build.outputs.image }
63     digest: ${ needs.build.outputs.digest }
64   steps:
65     - name: checkout repo
66     - name: Install cosign
67     - name: Install Syft
68     - name: Login to ghcr.io
69
70     - name: Attach SBOM to image
71       run: |
72         # syft pulls the image and analyses the contents to generate an SBOM
73         syft "${image}@${digest}" --output spdx-json --file sbom.spdx.json
74         cosign attest --predicate sbom-final.spdx.json --type spdxjson "${image}@${digest}" --yes
```


Pipeline - Generate & Attach SLSA Provenance

```
76 # slsa-github-generator creates and pushes the provenance attestation
77 provenance:
78   needs: [build]
79   permissions:
80     actions: read
81     id-token: write
82     # contents: read
83     packages: write
84   if: startsWith(github.ref, 'refs/tags/') # just to be safe, don't push if it's not a tag
85   uses: slsa-framework/slsa-github-generator/.github/workflows/generator_container_slsa3.yml@v1.9.0
86   with:
87     image: ${ needs.build.outputs.image }
88     digest: ${ needs.build.outputs.digest }
89     registry-username: ${ github.actor }
90     compile-generator: true
91   secrets:
92     registry-password: ${ secrets.GITHUB_TOKEN }
```

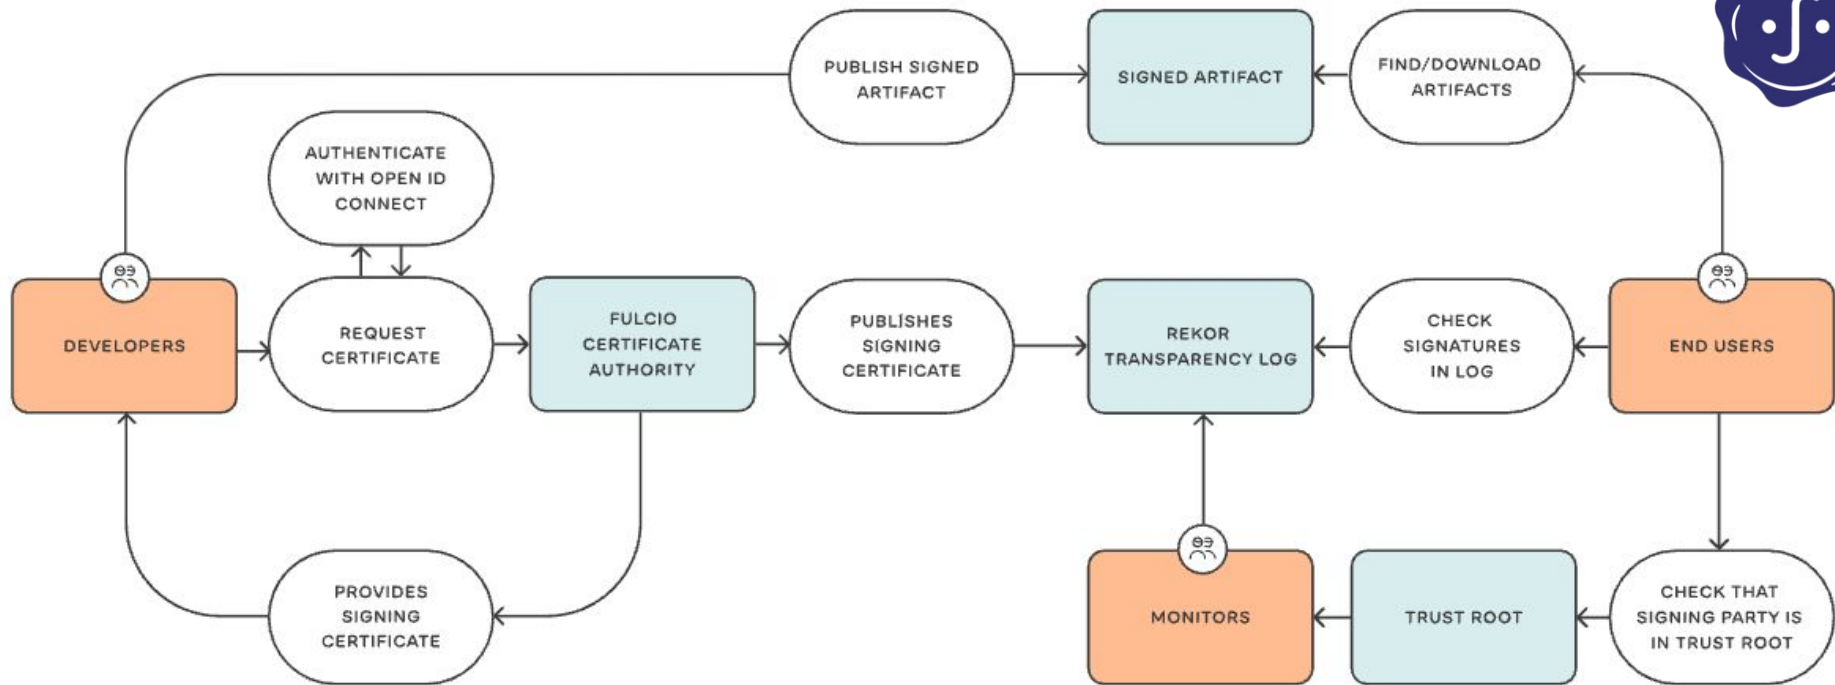

Open-source vs Private

Achieving SLSA (even Level 3) is fairly easy today for open-source projects, if you use the public sigstore instance and GitHub.

For private/proprietary projects, it's hard as you have to setup a code signing infrastructure. (also, how will consumers access it for verification?)

Signing

Sigstore



Verifications

- Image Signature
- SBOM
- SLSA Provenance

Verifying

Verifying the Image with Nerdctl / Finch

```
finch run -it \  
  --verify=cosign \  
  --cosign-certificate-identity=https://github.com/chainguard-images/images/.github/workflows/release.yaml@refs/heads/main \  
  --cosign-certificate-oidc-issuer=https://token.actions.githubusercontent.com \  
  cgr.dev/chainguard/busybox /bin/sh
```

Verifying and downloading the SBOM (+grype)

```
cosign verify-attestation ghcr.io/verifa/verinotes:v0.2.0 \
...--certificate-identity-regexp '^https://github.com/verifa/verinotes/.github/workflows/.*.yaml@refs/tags/v0.*.*' \
...--certificate-oidc-issuer https://token.actions.githubusercontent.com \
...--type spdxjson | jq '.payload | @base64d' -r | jq '.predicate' | grype -v
```



```
--type spdxjson | jq '.payload | @base64d' -r | jq '.predicate' | grype -v
[0000] INFO grype version: 0.64.1
[0000] INFO new version of grype is available: 0.73.3 (currently running: 0.64.1)
[0000] INFO downloading new vulnerability DB
```

Verification for ghcr.io/verifa/verinotes:v0.2.0 --

The following checks were performed on each of these signatures:

- The cosign claims were validated
- Existence of the claims in the transparency log was verified offline
- The code-signing certificate was verified using trusted certificate authority certificates

Certificate subject: <https://github.com/verifa/verinotes/.github/workflows/release.yaml@refs/tags/v0.2.0>

Certificate issuer URL: <https://token.actions.githubusercontent.com>

GitHub Workflow Trigger: push

GitHub Workflow SHA: d0a495e16d32c2bb125bf1481f747e711714199b

GitHub Workflow Name: Release workflow

GitHub Workflow Repository: verifa/verinotes

GitHub Workflow Ref: refs/tags/v0.2.0

```
[0001] WARN some package(s) are missing CPEs. This may result in missing vulnerabilities. You may autogenerate these using: --add-cpes-if-none
[0010] INFO updated vulnerability DB from version=5 built="2023-08-16 01:26:37 +0000 UTC" to version=5 built="2023-11-21 01:29:05 +0000 UTC"
[0010] INFO found 9 vulnerabilities for 300 packages
```

NAME	INSTALLED	FIXED-IN	TYPE	VULNERABILITY	SEVERITY
@sveltejs/kit	1.6.0	1.15.1	npm	GHSA-5p75-vc5g-8rv2	High
@sveltejs/kit	1.6.0	1.15.2	npm	GHSA-gv7g-x59x-wf8f	High
postcss	8.4.21	8.4.31	npm	GHSA-7fh5-64p2-3v2j	Medium
semver	7.3.8	7.5.2	npm	GHSA-c2qf-rxjj-qqgw	Medium
undici	5.18.0	5.19.1	npm	GHSA-5r9g-qh6m-jxff	Medium
undici	5.18.0	5.19.1	npm	GHSA-r6ch-mqf9-qc9w	High
undici	5.18.0	5.26.2	npm	GHSA-wqq4-5wpv-mx2g	Low
vite	4.1.1	4.1.5	npm	GHSA-353f-5xf4-qw67	High
word-wrap	1.2.3	1.2.4	npm	GHSA-j8xg-fqq3-53r7	Medium

Verifying SLSA provenance

```
cosign verify-attestation \  
  --type slsaprovenance \  
  --certificate-oidc-issuer https://token.actions.githubusercontent.com \  
  --certificate-identity-regexp '^https://github.com/slsa-framework/slsa-github-generator/.github/workflows/generator_container_slsa3.yml' \  
  --policy policy.cue \  
  ghcr.io/verifa/verinotes:v0.2.0 | jq '.payload' | @base64d' -r | jq '.predicate'
```

File: **policy.cue**

```
1 // The predicateType field must match this string
2 predicateType: "https://slsa.dev/provenance/v0.2"
3
4 predicate: {
5     // This condition verifies that the builder is the builder we
6     // expect and trust. The following condition can be used
7     // unmodified. It verifies that the builder is the container
8     // workflow.
9     builder: {
10        id: =~"^https://github.com/slsa-framework/slsa-github-generator/.github/workflows/gener
11    }
12    invocation: {
13        configSource: {
14            // This condition verifies the entrypoint of the workflow.
15            // Replace with the relative path to your workflow in your
16            // repository.
17            entryPoint: ".github/workflows/release.yaml"
18
19            // This condition verifies that the image was generated from
20            // the source repository we expect. Replace this with your
21            // repository.
22            uri: =~"^git\\+https://github.com/verifa/verinotes@refs/tags/v[0-9]+.[0-9]+.[0-9]+$"
23        }
24    }
25 }
```

```
cosign verify-attestation \  
  --type slsaprovenance \  
  --certificate-oidc-issuer https://token.actions.githubusercontent.com \  
  --certificate-identity-regexp '^https://github.com/slsa-framework/slsa-github-generator/.github/workflows/generator_container_slsa3.yml@refs\  
  --policy policy.cue \  
  ghcr.io/verifa/verinotes:v0.2.0 | jq '.payload | @base64d' -r | jq '.predicate'  
will be validating against CUE policies: [policy.cue]
```

Verification for ghcr.io/verifa/verinotes:v0.2.0 --

The following checks were performed on each of these signatures:

- The cosign claims were validated
- Existence of the claims in the transparency log was verified offline
- The code-signing certificate was verified using trusted certificate authority certificates

Certificate subject: https://github.com/slsa-framework/slsa-github-generator/.github/workflows/generator_container_slsa3.yml@refs/tags/v1.9.0

Certificate issuer URL: https://token.actions.githubusercontent.com

GitHub Workflow Trigger: push

GitHub Workflow SHA: d0a495e16d32c2bb125bf1481f747e711714199b

GitHub Workflow Name: Release workflow

GitHub Workflow Repository: verifa/verinotes

GitHub Workflow Ref: refs/tags/v0.2.0

```
{  
  "builder": {  
    "id": "https://github.com/slsa-framework/slsa-github-generator/.github/workflows/generator_container_slsa3.yml@refs/tags/v1.9.0"  
  },  
  "buildType": "https://github.com/slsa-framework/slsa-github-generator/container@v1",  
  "invocation": {  
    "configSource": {  
      "uri": "git+https://github.com/verifa/verinotes@refs/tags/v0.2.0",  
      "digest": {  
        "sha1": "d0a495e16d32c2bb125bf1481f747e711714199b"  
      },  
      "entryPoint": ".github/workflows/release.yaml"  
    },  
    "parameters": {},  
    "environment": {  
      "github_actor": "mvainio-verifa",
```

There is also an official
slsa-verifier

Kyverno / Policy Controller

Isn't there an easier way?

GOOD NEWS!

Platforms and open-source ecosystems are coming up with native integrations

GitHub

Actions Projects Wiki Security 10 Insights Settings

Dependency graph

Dependencies Dependents Dependabot

Export SBOM

Search all dependencies

@sveltejs/kit 1.6.0

2 high

Detected automatically on Aug 02, 2023 (npm) - ui/package-lock.json · MIT


npm

Provenance

Built and signed on
 **GitHub Actions**
[View build summary](#)

Source Commit github.com/prisma/prisma@bfe7bf8
Build File [.github/workflows/release-latest.yml](https://github.com/workflows/release-latest.yml)
Public Ledger [Transparency log entry](#)

I created this metadata,
now what?



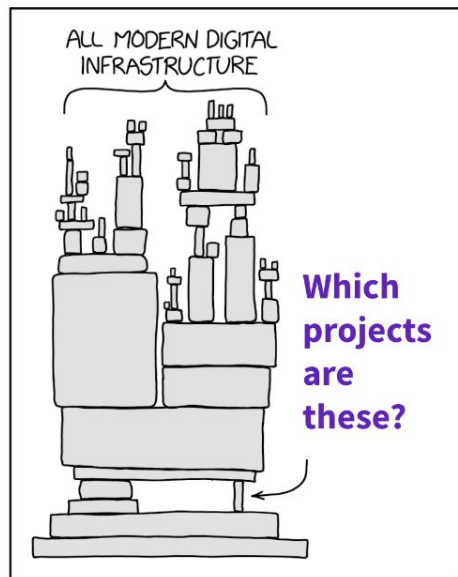
GUAC

Know your software supply chain

GUAC gives you directed, actionable insights into the security of your software supply chain.

Proactive

How do I prevent large scale supply chain compromises?



<https://xkcd.com/2347/>

Preventive

Have I taken the right safeguards?

When deciding to use and deploy software, are there sufficient security checks and approvals?



SLSA



trivy



grype

Reactive


HOW AM I AFFECTED???

A vulnerability or supply chain compromise is discovered!




Check it out: <https://github.com/guacsec/guac>

Public Record



SLSA, SBOM,
OSV, VEX, etc.

Org. (Private)



SLSA, SBOM,
Internal certification,
etc.


Vendors



trivy

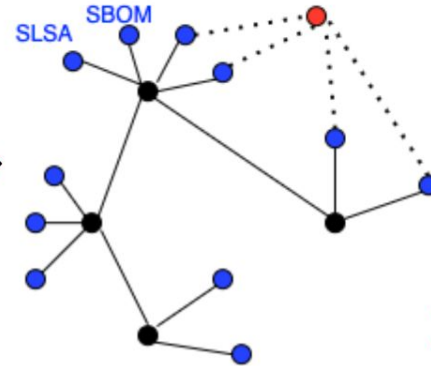


grype




Vuln certifications,
OSV, VEX, etc.

Extract
information and
relationships




Graph contains:

- Artifacts/identifiers
- Attestations
- Identities
- Relationships (edges)



CISO
Inventory
Policy
CMDB



Further reading

What is in-toto and how it relates to SLSA?

<https://slsa.dev/blog/2023/05/in-toto-and-slsa>

OpenPubKey vs sigstore (note: by sigstore maintainer):

<https://blog.sigstore.dev/openpubkey-and-sigstore/>

OpenVEX (& VEX in general):

<https://github.com/openvex>

Homebrew core going for SLSA Build level 2:

<https://blog.trailofbits.com/2023/11/06/adding-build-provenance-to-homebrew/>

CRI-O support for verifying image signatures:

<https://kubernetes.io/blog/2023/06/29/container-image-signature-verification/>

OpenSSF blog/github etc.

<https://openssf.org/blog/>