

Software Heritage

Towards A New Era for Software Engineering, Cybersecurity, and AI

Roberto Di Cosmo

Director, Software Heritage
Inria and Université Paris Cité

September 8–12, 2026
CBSOft 2026, São Paulo



Software Heritage
THE GREAT LIBRARY OF SOURCE CODE

- 1 Meet Software Heritage
- 2 From archive to revolutionary infrastructure
- 3 Software for research: reference archive ACAD_ADOPTION
- 4 Research on software: datasets from the archive DATASETS
- 5 Selected highlight: Impact on software studies ESE_STUDIES
- 6 A Very Large Telescope to explore the software development galaxy
- 7 Selected highlight: Improving Open Source Security with SWH(Sec) SECURITY
- 8 Selected highlight: AI and transparent LLMs CODECOMMONS
- 9 One last thing
- 10 Conclusion



Software Heritage
THE GREAT LIBRARY OF SOURCE CODE

Inria

with



unesco



The largest open source code archive: one infrastructure, open, shared, non profit

Unique digital common good built in France since 2015

Cultural Heritage



Source files

29,523,977,803

Industry



Commits

6,112,189,339


Research



Projects

441,196,029

Public Administration



5000+ platforms

All versions, all history
development in a single graph

- 58.7×10^9 nodes
- 1041×10^9 edges
~ 3 PB of storage

A revolutionary **infrastructure** ensures **availability** guarantees **integrity** enables **traceability**



HUAWEI



Infosys



Université Paris Cité



AdaCore



GitHub



Premier ministre



UNIVERSITÀ di PISA



SORBONNE UNIVERSITÉ



ÉCOLE NORMALE SUPÉRIEURE



UNIVERSITÉ DE LORRAINE



Saint'Anna



NWQ




université PARIS-SACLAY

The Software Hash identifier (SWHID)

Software Heritage Identifiers (SWHID)

see swhid.org



50+B
intrinsic,
decentralised,
cryptographic

Full fledged *source code references* for traceability, integrity and reproducibility

- Linux Foundation [SPDX 2.2](https://spdx.org/licenses/)
- IANA-registered "swh:"
- WikiData property [P6138](https://www.wikidata.org/wiki/P6138)

Examples: [Apollo 11 AGC excerpt](#), [Quake III rsqrt](#)
Guidelines available, see [the HOWTO](#)


[ISO/IEC 18670](#), see swhid.org

- 1 Meet Software Heritage
- 2 From archive to revolutionary infrastructure
- 3 Software for research: reference archive ACAD_ADOPTION
- 4 Research on software: datasets from the archive DATASETS
- 5 Selected highlight: Impact on software studies ESE_STUDIES
- 6 A Very Large Telescope to explore the software development galaxy
- 7 Selected highlight: Improving Open Source Security with SWH(Sec) SECURITY
- 8 Selected highlight: AI and transparent LLMs CODECOMMONS
- 9 One last thing
- 10 Conclusion


From archive to revolutionary infrastructure

Modern "Library of Alexandria", *international, non profit, long term* initiative addressing the needs of *industry, research, culture and society as a whole*


Software Graph




Software Blockchain



Open Science pillar



Big Code



Today we focus on

Open Science • **Large-scale analysis** • **Cybersecurity** • **AI**

- 1 Meet Software Heritage
- 2 From archive to revolutionary infrastructure
- 3 Software for research: reference archive ACAD_ADOPTION
- 4 Research on software: datasets from the archive DATASETS
- 5 Selected highlight: Impact on software studies ESE_STUDIES
- 6 A Very Large Telescope to explore the software development galaxy
- 7 Selected highlight: Improving Open Source Security with SWH(Sec) SECURITY
- 8 Selected highlight: AI and transparent LLMs CODECOMMONS
- 9 One last thing
- 10 Conclusion

An example is worth a thousand words

- Browse (e.g. [Apollo 11 \[excerpt\]](#), your work [may be already there !](#))
- Trigger archival, use the [updateswh](#) browser extension, configure the webhooks
- Get and use SWHIDs ([full specification available online](#))
- Cite software with [biblatex-software](#) package from CTAN
 - [Overleaf ACMART template](#) available
- Example in journals: [article from IPOL](#)
- Example with Parmap: [devel on Github](#), [archive in SWH](#), [curated deposit in HAL](#)
- Extracting all the software products [for Inria](#), [for CNRS](#), [for CNES](#), [for LIRMM](#) or [for Rémi Gribonval](#) using [HalTools](#)
- [Curated deposit in SWH via HAL](#), see for example: [LinBox](#), [SLALOM](#), [Givaro](#), [NS2DDV](#), [SumGra](#), [Coq proof](#), ...
- Example use in research articles:
 - compare Fig. 1 and conclusions in [the 2012 version](#) and [the updated version](#)
 - SWHID in [a replication experiment](#)

A few adoption indicators



Policy



- [Recommendations in ANR 2023 guidelines \(p. 17\)](#)
- HAL+SWH in [the Open Science software booklet](#)

Projects



FAIRCORE4EOSC
Core Components Supporting a FAIR EOSC

The CodeMeta Project



FAIR-IMPACT
Expanding FAIR solutions across EOSC

Users and collaborations

What are they “referencing”?

source	n	percentage
Not available	2868	46.22
GitHub	1151	18.55
software heritage	387	6.24
zenodo	142	2.29
r package	70	1.13
cran	56	0.90
r package version	54	0.87
gitlab	35	0.56

Graphics Replicability Stamp Initiative



b/Surf: Interactive Bézier Splines on Surface Meshes

Claudio Mancinelli, Giacomo Nazzaro, Fabio Pellacini, Enrico Puppo
IEEE Transactions on Visualization and Computer Graphics (TVCG)



Repository



Best practices to disseminate

Artifact Evaluation Committees

current status

- many guidelines developed in the 2010's ignore Software Heritage
- ACM DL or Zenodo deposit of a zip + DOI still required

state of the art

- archive on Software Heritage
- use the qualified SWHID

Software citation

state of the art

- include codemeta.json (or citation.cff, better if generated from codemeta.json)
- generate biblatex entry, and use biblatex-software (CTAN, acmart.cls, etc.)



Want the full story? The course notes ***The Source Code of Science***

Illustration: AI-generated (Google Gemini), author-directed; provenance record in the book's appendix.



- 1 Meet Software Heritage
- 2 From archive to revolutionary infrastructure
- 3 Software for research: reference archive ACAD_ADOPTION
- 4 Research on software: datasets from the archive DATASETS
- 5 Selected highlight: Impact on software studies ESE_STUDIES
- 6 A Very Large Telescope to explore the software development galaxy
- 7 Selected highlight: Improving Open Source Security with SWH(Sec) SECURITY
- 8 Selected highlight: AI and transparent LLMs CODECOMMONS
- 9 One last thing
- 10 Conclusion

<https://registry.opendata.aws/software-heritage/>

Registry of Open Data on AWS



Software Heritage Graph Dataset

digital preservation

free software

open source software

source code

Description

[Software Heritage](#) is the largest existing public archive of software source code and accompanying development history. The Software Heritage Graph Dataset is a fully deduplicated Merkle DAG representation of the Software Heritage archive. The dataset links together file content identifiers, source code directories, Version Control System (VCS) commits tracking evolution over time, up to the full states of VCS repositories as observed by Software Heritage during periodic crawls. The dataset's contents come from major development forges (including GitHub and GitLab), FOSS distributions (e.g., Debian), and language-specific package managers (e.g., PyPI). Crawling information is also included, providing timestamps about when and where all archived source code artifacts have been observed in the wild.

Update Frequency

Data is updated yearly

License

Creative Commons Attribution 4.0 International. By accessing the dataset, you agree with the Software Heritage [Ethical Charter for using the archive data](#) and the [terms of use for bulk access](#).

Documentation

<https://docs.softwareheritage.org/devel/swh-dataset/graph/athena.html>

Managed By

Software Heritage

See all datasets managed by [Software Heritage](#).

Contact

R. Di Cosmo

roberto@dicosmo.org

(CC-BY 4.0)

Resources on AWS

Description

Software Heritage Graph Dataset

Resource type

S3 Bucket

Amazon Resource Name (ARN)

```
arn:aws:s3:::softwareheritage
```

AWS Region

```
us-east-1
```

AWS CLI Access (No AWS account required)

```
aws s3 ls --no-sign-request s3://softwareheritage/
```

Description

[S3 Inventory](#) files

Resource type

S3 Bucket

Amazon Resource Name (ARN)

```
arn:aws:s3:::softwareheritage-inventory
```

AWS Region

```
us-east-1
```

AWS CLI Access (No AWS account required)

```
aws s3 ls --no-sign-request s3://softwareheritage-inventory
```

Software Heritage for SE, CyberSecurity and AI

September 2026

7 / 37

State-of-the-art graph compression from social networks



Paolo Boldi, Antoine Pietri, Sebastiano Vigna, Stefano Zacchiroli

Ultra-Large-Scale Repository Analysis via Graph Compression

SANER 2020, 27th Intl. Conf. on Software Analysis, Evolution and Reengineering. IEEE

Results

Full graph structure (58.7 B nodes, 1.04 T edges) in 400 GiB RAM

- traversal time is tens of ns per edge
- bidirectional traversals implemented
- **beware:** metadata access is still *off RAM*


Java Rust and gRPC APIs available ...

docs.softwareheritage.org/devel/swh-graph/grpc-api.html

- 1 Meet Software Heritage
- 2 From archive to revolutionary infrastructure
- 3 Software for research: reference archive ACAD_ADOPTION
- 4 Research on software: datasets from the archive DATASETS
- 5 Selected highlight: Impact on software studies ESE_STUDIES**
- 6 A Very Large Telescope to explore the software development galaxy
- 7 Selected highlight: Improving Open Source Security with SWH(Sec) SECURITY
- 8 Selected highlight: AI and transparent LLMs CODECOMMONS
- 9 One last thing
- 10 Conclusion

Mining Android Applications on Software Heritage


RQ: how to build a specific dataset for a given research question?



(from the Inria/IRISA DiverSE team)

Using the SWH merkle dag to identify android repositories

Identify android application repositories = Find the AndroidManifest.xml among the sources




SWH Merkle DAG, Antoine Pietri

1) Iterate over the graph nodes until you find a directory node containing a file named "AndroidManifest.xml".

Using the SWH merkle dag to identify android repositories

Identify android application repositories = Find the AndroidManifest.xml among the sources




SWH Merkle DAG, Antoine Pietri

2) Extract the SWH identifier of the blob corresponding to the AndroidManifest.xml and download the corresponding file through the SWH Web API

Using the SWH merkle dag to identify android repositories

Identify android application repositories = Find the AndroidManifest.xml among the sources



SWH Merkle DAG, Antoine Pietri

3) Traverse the graph in backward direction to the origin node and get the repository url

Broad variety of sources in *one open dataset*

reduces usual GH bias

Reference simple *standard data format*

VCS and forge details are abstracted away

Simplifies reproducibility packages

no need to create a full copy, *just list the SWHIDs!*

Software Heritage does the heavy lifting for you


no need to scrape/download repositories all over again

Example: (Open) Source Code comes from all over the world

MSR '22, May 23–24, 2022, Pittsburgh, PA, USA

Davide Rossi and Stefano Zacchioli

Geographic Diversity in Public Code Contributions: An Exploratory Large-Scale Study Over 50 Years. MSR 2022) <https://doi.org/10.1145/3524842.3528471>




We use as dataset the Software Heritage archive [3] and analyze from it 2.2 billion commits archived from 160 million projects and authored by 43 million authors during the 1971–2021 time period. We geolocate developers to 12 world regions, using as signals email country code top-level domains (ccTLDs) and author (first/last) names compared with name distributions around the world, and UTC offsets mined from commit metadata.



Software Heritage
THE GREAT LIBRARY OF SOURCE CODE

Figure 3: Ratio of commits (above) and active authors (below) by world zone over the 1971–2020 period.

Example: Programming language evolution in (Open) Source



A. Desmazières, R. Di Cosmo, V. Lorentz

50 years of programming language evolution through the Software Heritage Looking Glass

International Conference on Mining Software Repositories (MSR) 2025.

<https://hal.science/hal-04924849/>

R. Di Cosmo

roberto@dicosmo.org

(CC-BY 4.0)


Software Heritage for SE, CyberSecurity and AI

September 2026

16 / 37



Example: (Open) Source Code grows at an exponential rate



Get the data
and the code



Example: Institutional contribution analysis

Question:


can we leverage the Software Heritage graph to map contributions from a research institution in the galaxy of software development?

Answer: let's try

- **Extract:** projects and contributions using email domain matching in the Software Heritage graph.
- **Deduplicate:** Cluster forks/copies using commit metadata, select a canonical repository.
- **Enrich:** Add GitHub metadata (topics, README, etc.) for context and visibility.
- **Analyze:** Compute key metrics on impact, contributor domains, and research relevance.

- 1 Meet Software Heritage
- 2 From archive to revolutionary infrastructure
- 3 Software for research: reference archive ACAD_ADOPTION
- 4 Research on software: datasets from the archive DATASETS
- 5 Selected highlight: Impact on software studies ESE_STUDIES
- 6 A Very Large Telescope to explore the software development galaxy**
- 7 Selected highlight: Improving Open Source Security with SWH(Sec) SECURITY
- 8 Selected highlight: AI and transparent LLMs CODECOMMONS
- 9 One last thing
- 10 Conclusion

From the Software Sky to Institutional Insight: Anatomy of an Observational Pipeline



**1445630****Contributions****37027****Software Projects****Active last 12 months: 489****Still crawled: 33269****14040****Contributors****1993-07-28**
2026-03-05**Time Span**

38 BR academic domains, 10+ contributions with a BR academic email:


USP, UNICAMP, UNESP, UFRJ, UFMG, UFRGS, UnB, IMPA, LNCC, CBPF, FIOCRUZ, EMBRAPA, INPE, PUC-Rio, FGV, FAPESP, CNPq, ...


1445630
 Contributions


37027
 Software Projects
 Active last 12 months: 489
 Still crawled: 33269



14040
 Contributors


1993-07-28
2026-03-05
 Time Span



domain	count
unibz.br	2973
ufpe.br	2103
unicamp.br	1236
uffir.br	1124
ufsj.br	1108
ufc.br	806
puccs.br	785
ufpa.br	436
ufmg.br	415
ufpi.br	376
unesp.br	363
ufem.br	317
ufac.br	277
puccrj.br	273
unb.br	247
ufu.br	229
ufscar.br	190
uffa.br	154
unifesp.br	97
ufpb.br	75
fgv.br	59
fla.br	39
imp.br	36
embrapa.br	33
faccru.br	33
ueff.br	33
lgt.br	18
incc.br	14
puccp.br	8
crisp.br	6
cbpt.br	4
fapesp.br	2

Software Heritage is **the only infrastructure** capable of producing these metrics



A Software Heritage product

The State of Public Code

How are governments contributing to Open Source?

[Read the report 1.1](#)

[Get the data](#)

[Get involved](#)

[About](#)



288,411

Repositories with public sector contributions



92,831

Public sector contributors with at least 1 commit



112/193

UN Member States with contributions to more than 10 repos

Source: Software Heritage archive snapshot, March 2026

- **Brazil is second worldwide:** 17,725 public sector contributors (US 41,253, UK 9,252)
- gov.br is the **largest government email domain** on the planet, nearly 2x the next one
- Brazilian contributors reach **~4.4 projects each** — world median ~2.2



stateofpubliccode.org

With UN ODET and UNESCO • gov-domain contributions: a floor, not a ceiling

- 1 Meet Software Heritage
- 2 From archive to revolutionary infrastructure
- 3 Software for research: reference archive ACAD_ADOPTION
- 4 Research on software: datasets from the archive DATASETS
- 5 Selected highlight: Impact on software studies ESE_STUDIES
- 6 A Very Large Telescope to explore the software development galaxy
- 7 Selected highlight: Improving Open Source Security with SWH(Sec) SECURITY**
- 8 Selected highlight: AI and transparent LLMs CODECOMMONS
- 9 One last thing
- 10 Conclusion

Programme

France 2030 PTCC — 5 years (2024–2028) — 8 research teams

Co-leaders: Olivier Barais (Inria, DiverSE) · Stefano Zacchiroli (Télécom Paris, ACES)

Partners — academic

- **APR** — LIP6, **Sorbonne Université** (*prog. languages, static analysis MOPSA, verification*)
- **ACES** — Télécom Paris (*cybersecurity, software supply chain*)
- **CEA List** (*static analysis: Frama-C, Binsec*)


Partners — Inria

- **DiverSE** — supply chain security
- **RMOD** — reverse engineering (Moose)
- **Software Heritage** — archive, infrastructure
- **Spirals** — self-protective systems, web security
- **Whisper** — Linux kernel, Coccinelle

Software supply chain attacks


Reusing OSS via dependencies

- **Software dependencies**: a popular way of reusing open source software.
- Software product *A* uses functionalities implemented in OSS product *B* ... and so on.




Attacking the software supply chain

- Attacking **undermaintained "leaf" packages** (e.g., D) → efficient attack strategy
- Many documented attacks: event-stream (2018), node-ipc (2022), XZ utils (2024), ...



based on xkcd.com/2347

- Any change to a piece of software (*commit*) can **introduce a new vulnerability**.
- Or it can **fix an existing vulnerability**.
- What happens if a project is forked **between introduction and fix** of a vulnerability?
- It inherits the vulnerability, ... until the change with the fix is integrated.



An universal knowledge base about public code vulnerabilities

Vision

- **Software Heritage** is the perfect (and only) place where to build an universal knowledge base that maps known vulnerabilities to public code artifacts.
- SWH can provide an **open data API mapping SWHIDs to CVEs**, that knows about *all public commits* and can be leveraged to increase open source security.

EU Cyber Resilience Act (CRA)

- Key helper to abide to CRA obligations, coming into effect ~Q3 2026.
- SWH funding member of the **Open Regulatory Compliance Working Group**.



Roadmap

- Context: SWHSec project (PTCC-funded, 2023-2027).
- Current status: working prototype that processes OSV.dev data and use it to "color" the entire SWH commit graph (~6 billion commits) with vulnerability information.



Syful Islam, Stefano Zacchioli.

On the Informativeness of Security Commit Messages: A Large-scale Replication Study.

To appear in EASE 2026, ACM, Glasgow, UK.



Syful Islam, Stefano Zacchioli.

On the Use of Commit Messages for Corrective Software Maintenance: A Systematic Mapping Study.

To appear in EASE 2026, ACM, Glasgow, UK.



Romain Robbes, Théo Matricon, Thomas Degueule, André Hora, Stefano Zacchioli.

Promises, Perils, and (Timely) Heuristics for Mining Coding Agent Activity.

To appear in MSR 2026, ACM, Rio de Janeiro, Brazil.



Luís Soeiro, Thomas Robert, Stefano Zacchioli.

Finding Software Supply Chain Attack Paths with Logical Attack Graphs.

FPS 2025, Brest, France. Springer, 2026.



Solal Rapaport, Laurent Pautet, Samuel Tardieu, Stefano Zacchioli.

Altered Histories in Version Control System Repositories: Evidence from the Trenches.

ASE 2025, IEEE, Seoul, South Korea. Pages 2183–2194.

- 1 Meet Software Heritage
- 2 From archive to revolutionary infrastructure
- 3 Software for research: reference archive ACAD_ADOPTION
- 4 Research on software: datasets from the archive DATASETS
- 5 Selected highlight: Impact on software studies ESE_STUDIES
- 6 A Very Large Telescope to explore the software development galaxy
- 7 Selected highlight: Improving Open Source Security with SWH(Sec) SECURITY
- 8 Selected highlight: AI and transparent LLMs CODECOMMONS**
- 9 One last thing
- 10 Conclusion

Software Heritage and Generative AI, first contacts

October 19, 2023

Software Heritage Statement on Large Language Models for Code



Principles

1. Knowledge derived from the Software Heritage archive must be given back to humanity, rather than monopolized for private gain. The resulting machine learning models must be made available under a suitable open license, together with the documentation and toolings needed to use them.
2. The initial training data extracted from the Software Heritage archive must be fully and precisely identified by, for example, publishing the corresponding SWHID identifiers (note that, in the context of Software Heritage, public availability of the *initial training data* is a given: anyone can obtain it from the archive). This will enable use cases such as: studying biases (fairness), verifying if a code of interest was present in the training data (transparency), and providing appropriate attribution when generated code bears resemblance to training data (credit), among others.
3. Mechanisms should be established, where possible, for authors to exclude their archived code from the training inputs before model training begins.

Software Heritage and Generative AI, first contacts

October 19, 2023

Software Heritage Statement on Large Language Models for Code





Principles

1. Knowledge derived from the Software Heritage archive must be given back to humanity, rather than monopolized for private gain. The resulting machine learning models must be made available under a suitable open license, together with the documentation and toolings needed to use them.
2. The initial training data extracted from the Software Heritage archive must be fully and precisely identified by, for example, publishing the corresponding SWHID identifiers (note that, in the context of Software Heritage, public availability of the initial training data is a given: anyone can obtain it from the archive). This will enable use cases such as: studying biases (fairness), verifying if a code of interest was present in the training data (transparency), and providing appropriate attribution when generated code bears resemblance to training data (credit), among others.
3. Mechanisms should be established, where possible, for authors to exclude their archived code from the training inputs before model training begins.

February 2024

Yes, it's possible!



Software Heritage and Generative AI, first contacts

October 19, 2023

Software Heritage Statement on Large Language Models for Code





Principles

1. Knowledge derived from the Software Heritage archive must be given back to humanity, rather than monopolized for private gain. The resulting machine learning models must be made available under a suitable open license, together with the documentation and toolings needed to use them.
2. The initial training data extracted from the Software Heritage archive must be fully and precisely identified by, for example, publishing the corresponding SWHID identifiers (note that, in the context of Software Heritage, public availability of the initial training data is a given: anyone can obtain it from the archive). This will enable use cases such as: studying biases (fairness), verifying if a code of interest was present in the training data (transparency), and providing appropriate attribution when generated code bears resemblance to training data (credit), among others.
3. Mechanisms should be established, where possible, for authors to exclude their archived code from the training inputs before model training begins.

February 2024

Yes, it's possible!



But it's hard...

Lessons learned

Principles

1. Knowledge derived from the Software Heritage archive must be given back to humanity, rather than monopolized for private gain. The resulting *machine learning models* must be made available under a suitable open license, together with the documentation and toolings needed to use them.
2. The *initial training data* extracted from the Software Heritage archive must be fully and precisely identified by, for example, publishing the corresponding **SWHID identifiers** (note that, in the context of Software Heritage, public availability of the *initial training data* is a given: anyone can obtain it from the archive). This will enable use cases such as: studying biases (fairness), verifying if a code of interest was present in the training data (transparency), and providing appropriate attribution when generated code bears resemblance to training data (credit), among others.
3. Mechanisms should be established, where possible, for authors to exclude their archived code from the training inputs before model training begins.

Lessons learned

Principles

1. Knowledge derived from the Software Heritage archive must be given back to humanity, rather than monopolized for private gain. The resulting *machine learning models* must be made available under a suitable open license, together with the documentation and toolings needed to use them.
2. The *initial training data* extracted from the Software Heritage archive must be fully and precisely identified by, for example, publishing the corresponding SWHID identifiers (note that, in the context of Software Heritage, public availability of the *initial training data* is a given: anyone can obtain it from the archive). This will enable use cases such as: studying biases (fairness), verifying if a code of interest was present in the training data (transparency), and providing appropriate attribution when generated code bears resemblance to training data (credit), among others.
3. Mechanisms should be established, where possible, for authors to exclude their archived code from the training inputs before model training begins.



Transparency is easy: [SWHID](#) (undergoing ISO standardisation) and Software Heritage
N.B. : may be mandated by regulations!

Opt out is complex: who is *the real right owner*?
(similar issues to license compliance)

- **Building the training set** is complex: e.g. includes **license compliance** alike work **at massive scale**
- Generating **attribution information** on model output **is more complex** than license compliance

We need a **coordinated effort** to ensure fully open models will succeed!

“Good enough” is not good enough *why the open web is collected so badly*

Good enough is **not** good enough

Why the open web is collected so
badly, and what it is costing us.

Roberto Di Cosmo · dicosmo.org/good-enough




dicosmo.org/good-enough


- **Massive multiple scraping** — every collector re-fetches every site, over and over, for material that has not changed
- **AI bots** — Software Heritage itself was hit by waves of automated collection: services unresponsive, engineers blocking addresses instead of building
- **Waste of resources** — bills explode for everyone, and the open web closes down one emergency measure at a time

CodeCommons: collect once, share widely — *one logical collection, many copies*

A STEP FORWARD: CodeCommons



CodeCommons: bird's eye view (technical focus)



CodeCommons


Open, responsible, and transparent AI: Our shared goal

CodeCommons is an ambitious project to create the world's most comprehensive digital commons for code


Building on the existing foundation of Software Heritage, the largest publicly available source code archive, CodeCommons aims to bring into one place all the **critical** and **qualified** information needed to create **smaller, better** datasets for the next generation of AI tools.

At its core, the project prioritizes transparency and traceability, enabling model builders and users to **respect creators' rights** while promoting **sovereign** and **sustainable** AI.

Learn more




Meet the teams



- 1 Meet Software Heritage
- 2 From archive to revolutionary infrastructure
- 3 Software for research: reference archive ACAD_ADOPTION
- 4 Research on software: datasets from the archive DATASETS
- 5 Selected highlight: Impact on software studies ESE_STUDIES
- 6 A Very Large Telescope to explore the software development galaxy
- 7 Selected highlight: Improving Open Source Security with SWH(Sec) SECURITY
- 8 Selected highlight: AI and transparent LLMs CODECOMMONS
- 9 One last thing**
- 10 Conclusion

Deep (infra)structural dependencies!


An early warning: left-pad (March 2016)



See the [Wikipedia](#) article:

- Facebook, PayPal, Netflix, Spotify affected
- web broken worldwide ~2.5 hours

A clear and present danger



Used by **Brazilian academic** devs (Source: *Software Heritage*)


"If GitHub or PyPI disappeared tomorrow, research would not just slow down — it would stop."

An ecumenical approach to sovereign fallback of build chains is possible

Add support for fallback in key package managers

npm, PyPI, Maven Central, Cargo, CPAN, RubyGems, Packagist, ...

Realize ecumenical resilience



“When a registry fails — the world continues to build.”

Four sovereign mirrors today.

Brazil could host the first one beyond Europe.

- 1 Meet Software Heritage
- 2 From archive to revolutionary infrastructure
- 3 Software for research: reference archive ACAD_ADOPTION
- 4 Research on software: datasets from the archive DATASETS
- 5 Selected highlight: Impact on software studies ESE_STUDIES
- 6 A Very Large Telescope to explore the software development galaxy
- 7 Selected highlight: Improving Open Source Security with SWH(Sec) SECURITY
- 8 Selected highlight: AI and transparent LLMs CODECOMMONS
- 9 One last thing
- 10 Conclusion

A growing and active community


Core Team



All together, 2026 Symposium



Ambassadors



A rally flag for a grand vision

Bring together academia, industry, governments, communities

"to build a reference, global infrastructure for open and better software"

Software Heritage is

- vendor neutral, open source
- a worldwide, long term initiative

Call to action

- leverage for AI, Cyber, SwEng
- get involved research, code

Software Heritage enables

- archival, reference, integrity
- qualification, sharing and reuse

- promote for AEC, Journals, etc.
- fund the long term mission

Annual report 2025



5 years in 5 minutes



Get these slides

