Poster: Building Networks of Trust for Legal Preservation with TAF

Dušan Nikolić
University of Novi Sad
nikolic.dusan@uns.ac.rs

Renata Vaderna
Independent Researcher
vrenata8@gmail.com

Patrick Zielinski
New York University
patrick.z@nyu.edu

David Greisen
Open Law Library
dgreisen@openlawlib.org

BJ Ard
University of Wisconsin–Madison
bj.ard@wisc.edu

Justin Cappos
New York University
jcappos@nyu.edu

Abstract: Digital law repositories require security and preservation guarantees surpassing those needed for long-term archival of other types of content and need to extend beyond any single publisher's lifespan. The Archive Framework (TAF) integrates Git version control with TUF to create tamper-evident legal repositories and lays the foundation for networks of trustworthy institutions to independently mirror publishers' repositories. Such networks address vulnerabilities no single publisher can mitigate alone, including publisher abandonment and coordinated pressure to alter history. However, deploying TAF and building these networks in practice introduces distinct operational challenges.

We present the first institutional deployment at the National Indian Law Library (NILL), which mirrors and validates legal repositories from five tribal governments through fully automated pipelines, demonstrating that institutional preservation networks can operate with minimal infrastructure.

I. Introduction

Legal preservation is an especially challenging domain: attacks against digital law repositories can have severe consequences, and preservation requirements extend decades or even centuries into the future. The Archive Framework (TAF) [1] combines Git version control with The Update Framework (TUF) [2] to create tamper-evident digital law repositories in which every version is cryptographically verifiable, enabling detection of unauthorized changes over time.

However, even repositories secured with TAF remain vulnerable to total infrastructure compromise or publisher abandonment. To mitigate these risks, TAF allows other entities to mirror law repositories. Building upon this foundation, this work focuses on the design and deployment of a network of independent, trustworthy institutions, such as law libraries and archives, together with network-level challenges not addressed by TAF alone.

Building such a network requires identifying and engaging participants and keeping setup and ongoing operation simple enough to function without specialized IT staff or dedicated hardware. The solution should integrate easily into existing institutional workflows and infrastructure. Participating institutions must keep replicated copies synchronized with minimal delay and verify all updates. In addition, the network needs to support detection of and recovery from catastrophic compromises, including scenarios in which a publisher or multiple institutions are compromised simultaneously.

The National Indian Law Library (NILL) [3] operates the first institutional implementation of this model, mirroring repositories from five tribal governments through fully automated pipelines and providing initial operational experience.

II. Background

TAF provides a foundation for creating tamper-evident legal archives and serves as the baseline system in this work. A TAF-based legal archive consists of target repositories storing legal materials and an authentication repository containing signed metadata that describes the expected state of each target over time. Attempts to remove historical commits or push updates without meeting required signing thresholds break the cryptographic verification chain. TAF's dependency mechanism allows authentication repositories to reference and validate other authentication repositories, enabling relationships between multiple TAF archives.

TAF includes an updater component that securely retrieves updates, verifies their authenticity, and rejects invalid states. The updater also supports post-update hooks that can be used to integrate validated updates into custom workflows.

TAF is currently deployed in production by fourteen jurisdictions in the United States, including Washington, D.C., the State of Maryland, and the City of Baltimore.

III. Network Design

The network operates on the principle that institutions witness rather than vote. Participants may mirror repositories from publishers or other institutions, and by doing so assert trust in the legitimacy of the mirrored party. This trust is established through bootstrapping, typically via direct communication that provides a known-good initial state, from which subsequent updates are validated automatically. Publishers may enable broader mirroring through mechanisms such as DNS-based authentication or restrict it to explicitly approved institutions.

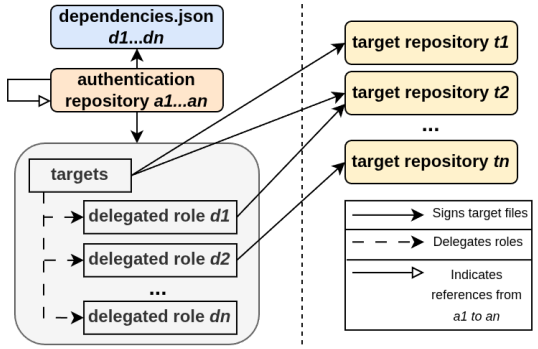

Fig. 1: Cross-repository trust in TAF. Tribal authentication repositories ($a_1 \dots a_n$) are listed in `dependencies.json`, a TAF configuration file, and NILL's top-level `targets` role delegates to one role per target repository ($d_1 \dots d_n$)

The network does not require a minimum number of participants.

Relationships between network members are recorded using TAF's dependency mechanism, as shown in Figure 1. Each institution maintains its own authentication repository and explicitly lists the authentication repositories it mirrors. For each mirrored repository, the listing specifies an initial state obtained through out-of-band authentication, along with one or more authoritative locations from which the repository can be cloned. This information enables the TAF updater to clone, validate, and subsequently update mirrored repositories.

As updates are retrieved and validated, institutions record the observed repository states as signed records within their own authentication repositories. Because these records are append-only and cryptographically bound to specific points in time, they form a transparency log. At the network level, independently maintained logs allow institutions to compare historical views of a repository and detect inconsistencies, for example when a publisher later denies or alters past content. While this approach supports recovery from isolated failures or compromised publishers, resolving discrepancies across logs remains an open problem and an important area for future work. In the legal domain, the possibility of coordinated compromise affecting multiple institutions means that simple consensus or majority-based approaches are insufficient.

Finally, broad participation requires lowering barriers to entry for institutions without dedicated IT staff or specialized infrastructure. To address this, institutions initialize their repositories and register signing keys using three JSON configuration files and a single CLI command, requiring only Python and the TAF package. To ease integration into existing workflows and automate tasks such as updating public-facing websites after successful repository updates, institutions may optionally use post-update hooks, as discussed in the following section.

IV. NILL DEPLOYMENT

NILL operates as the first library in this networked preservation model, serving an important mission: enhancing the power of tribal courts and strengthening tribal sovereignty by providing tribal leaders, legal practitioners, and the public with convenient access to current and accurate copies of tribal codes and constitutions. One aspect of tribal sovereignty is ensuring that Tribal Nations remain in control of their own information and how that information is used. When a Tribal Nation grants permission to share its legal resources in NILL's collection, that permission does not extend to any other use, and NILL cannot consent on behalf of Tribal Nations for inclusion in other databases or collections.

NILL mirrors legal repositories from five tribal governments, with validation and publication automated through GitHub Actions. Every 24 hours, a scheduled workflow triggers the TAF updater, which validates new states and updates NILL's copies of all referenced repositories, recording the update results.

Once the updater completes successfully, TAF's pipeline handler executes custom scripts tailored to NILL's publication requirements. These scripts transform the raw HTML content from each tribe's repository into a format suitable for the Tribal Law Gateway [4]. Specifically, the scripts apply NILL's standardized CSS styling, update URL references to point to the appropriate Gateway endpoints, and inject consistent header and footer elements across all pages. The result is a set of library-ready HTML files that preserve each tribe's legal content while presenting it within NILL's unified interface.

These transformed files are pushed to a separate Git repository, from which NILL's IT administrators can deploy updates to the static hosting infrastructure at their discretion. This separation of concerns allows the automated validation and transformation pipeline to run continuously, while leaving final control over public-facing changes with NILL staff.

V. CONCLUSION AND CHALLENGES

We have presented an approach to more robust long-term preservation of legal data through a network of trustworthy institutions that mirror legal repositories and record signed evidence of upstream changes. The NILL deployment, serving five tribal governments, demonstrates that institutions can participate in such a network without requiring specialized IT staff, while increasing resilience against failures or compromise compared to relying on a single archive alone.

Key challenges for future work include resolving disputes when institutional logs disagree, growing and sustaining the network as institutions join and leave amid funding and organizational changes, and bootstrapping initial trust using mechanisms beyond DNS-based authentication.

REFERENCES

[1] R. Vaderna, D. Nikolic, P. Zielinski, D. Greisen, B. Ard, and J. Cappos, "Enhancing legal document security and accessibility with taf," *Univ. of Wisconsin Legal Studies Research Paper Forthcoming*, 2025.

[2] The Update Framework (TUF), "The update framework (tuf)," https://theupdateframework.io, 2026, accessed: 2026-01-19.

[3] National Indian Law Library (NILL), "National indian law library (nill)," https://narf.org/nill, 2026, accessed: 2026-01-19.

[4] Tribal Law Gateway, "Tribal law gateway," https://narf.org/nill/triballaw/index.html, 2026, accessed: 2026-01-19.

Building Networks of Trust for Legal Preservation: The NILL Deployment

Dušan Nikolić[1], Renata Vaderna, Patrick Zielinski[2], David Greisen[3], BJ Ard[4], Justin Cappos[2]

[1] [2] [3] [4]

Concept

Who preserves authenticatable law when publishers disappear or face pressure to rewrite history?

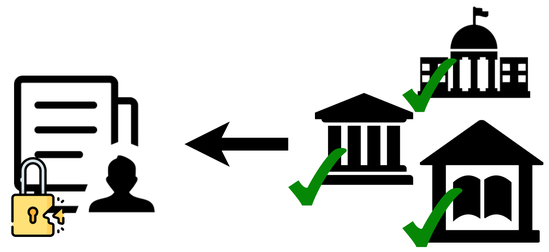

Institutions witness, don't vote

Solution: Independent institutions mirror repositories and maintain signed transparency logs of all updates.

Deployment with NILL

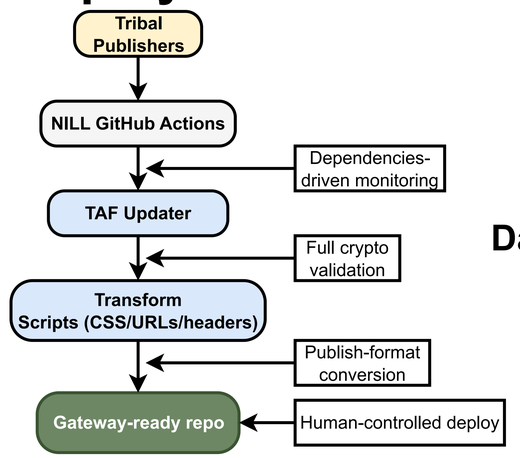

5 publishers mirrored

Daily end-to-end validation

No dedicated hardware

TAF Overview

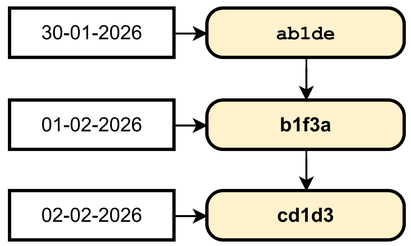

git TUF

A TAF archive consists of target Git repos (legal data) and an authentication repo (signed TUF metadata kept in Git).

Clients verify updates against signed metadata, detecting tampering.

Authentication Repository

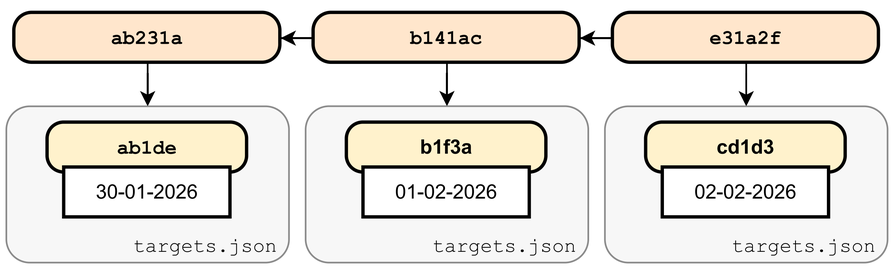

Networks of Trust With TAF

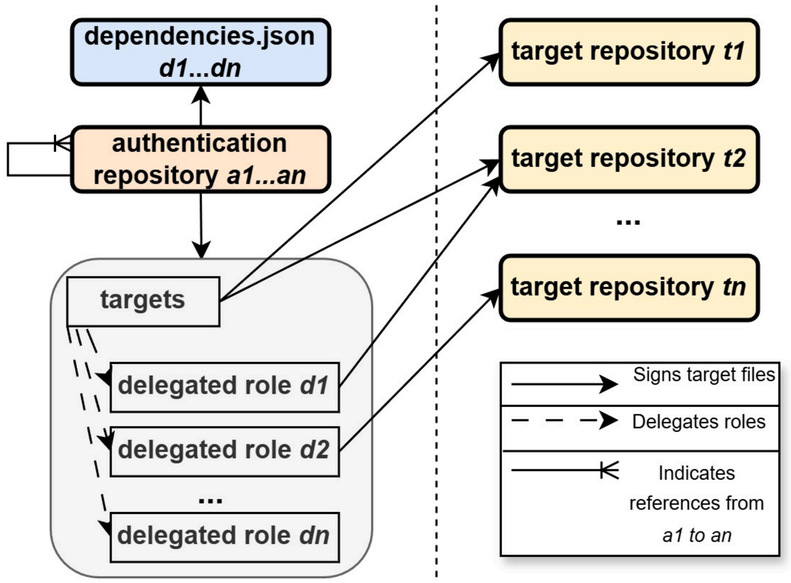

To scale beyond single publishers, institutions declare who they trust and mirror in **dependencies.json**—a signed configuration file listing other authentication repositories to mirror. This enables auditable, automated mirroring.

As institutions validate and update dependencies, they record observed repository states as signed target files, forming transparency logs. These logs provide cryptographic evidence of what each institution witnessed.

Implementation

NILL

TAF

Conclusion

Takeaway: Independent institutions can verify and log legal history, not just host it.

Next: disagreement handling, scalable trust bootstrapping, long-term sustainability.

Threat	Network evidence	Open challenge
Publisher compromise / rewrite	Independent signed logs	Multiple independent mirrors
Publisher disappears	Mirrors keep serving	Long-term sustainability (funding/ops)
Coordinated pressure	Mirrors as witnesses	Dispute resolution when logs disagree