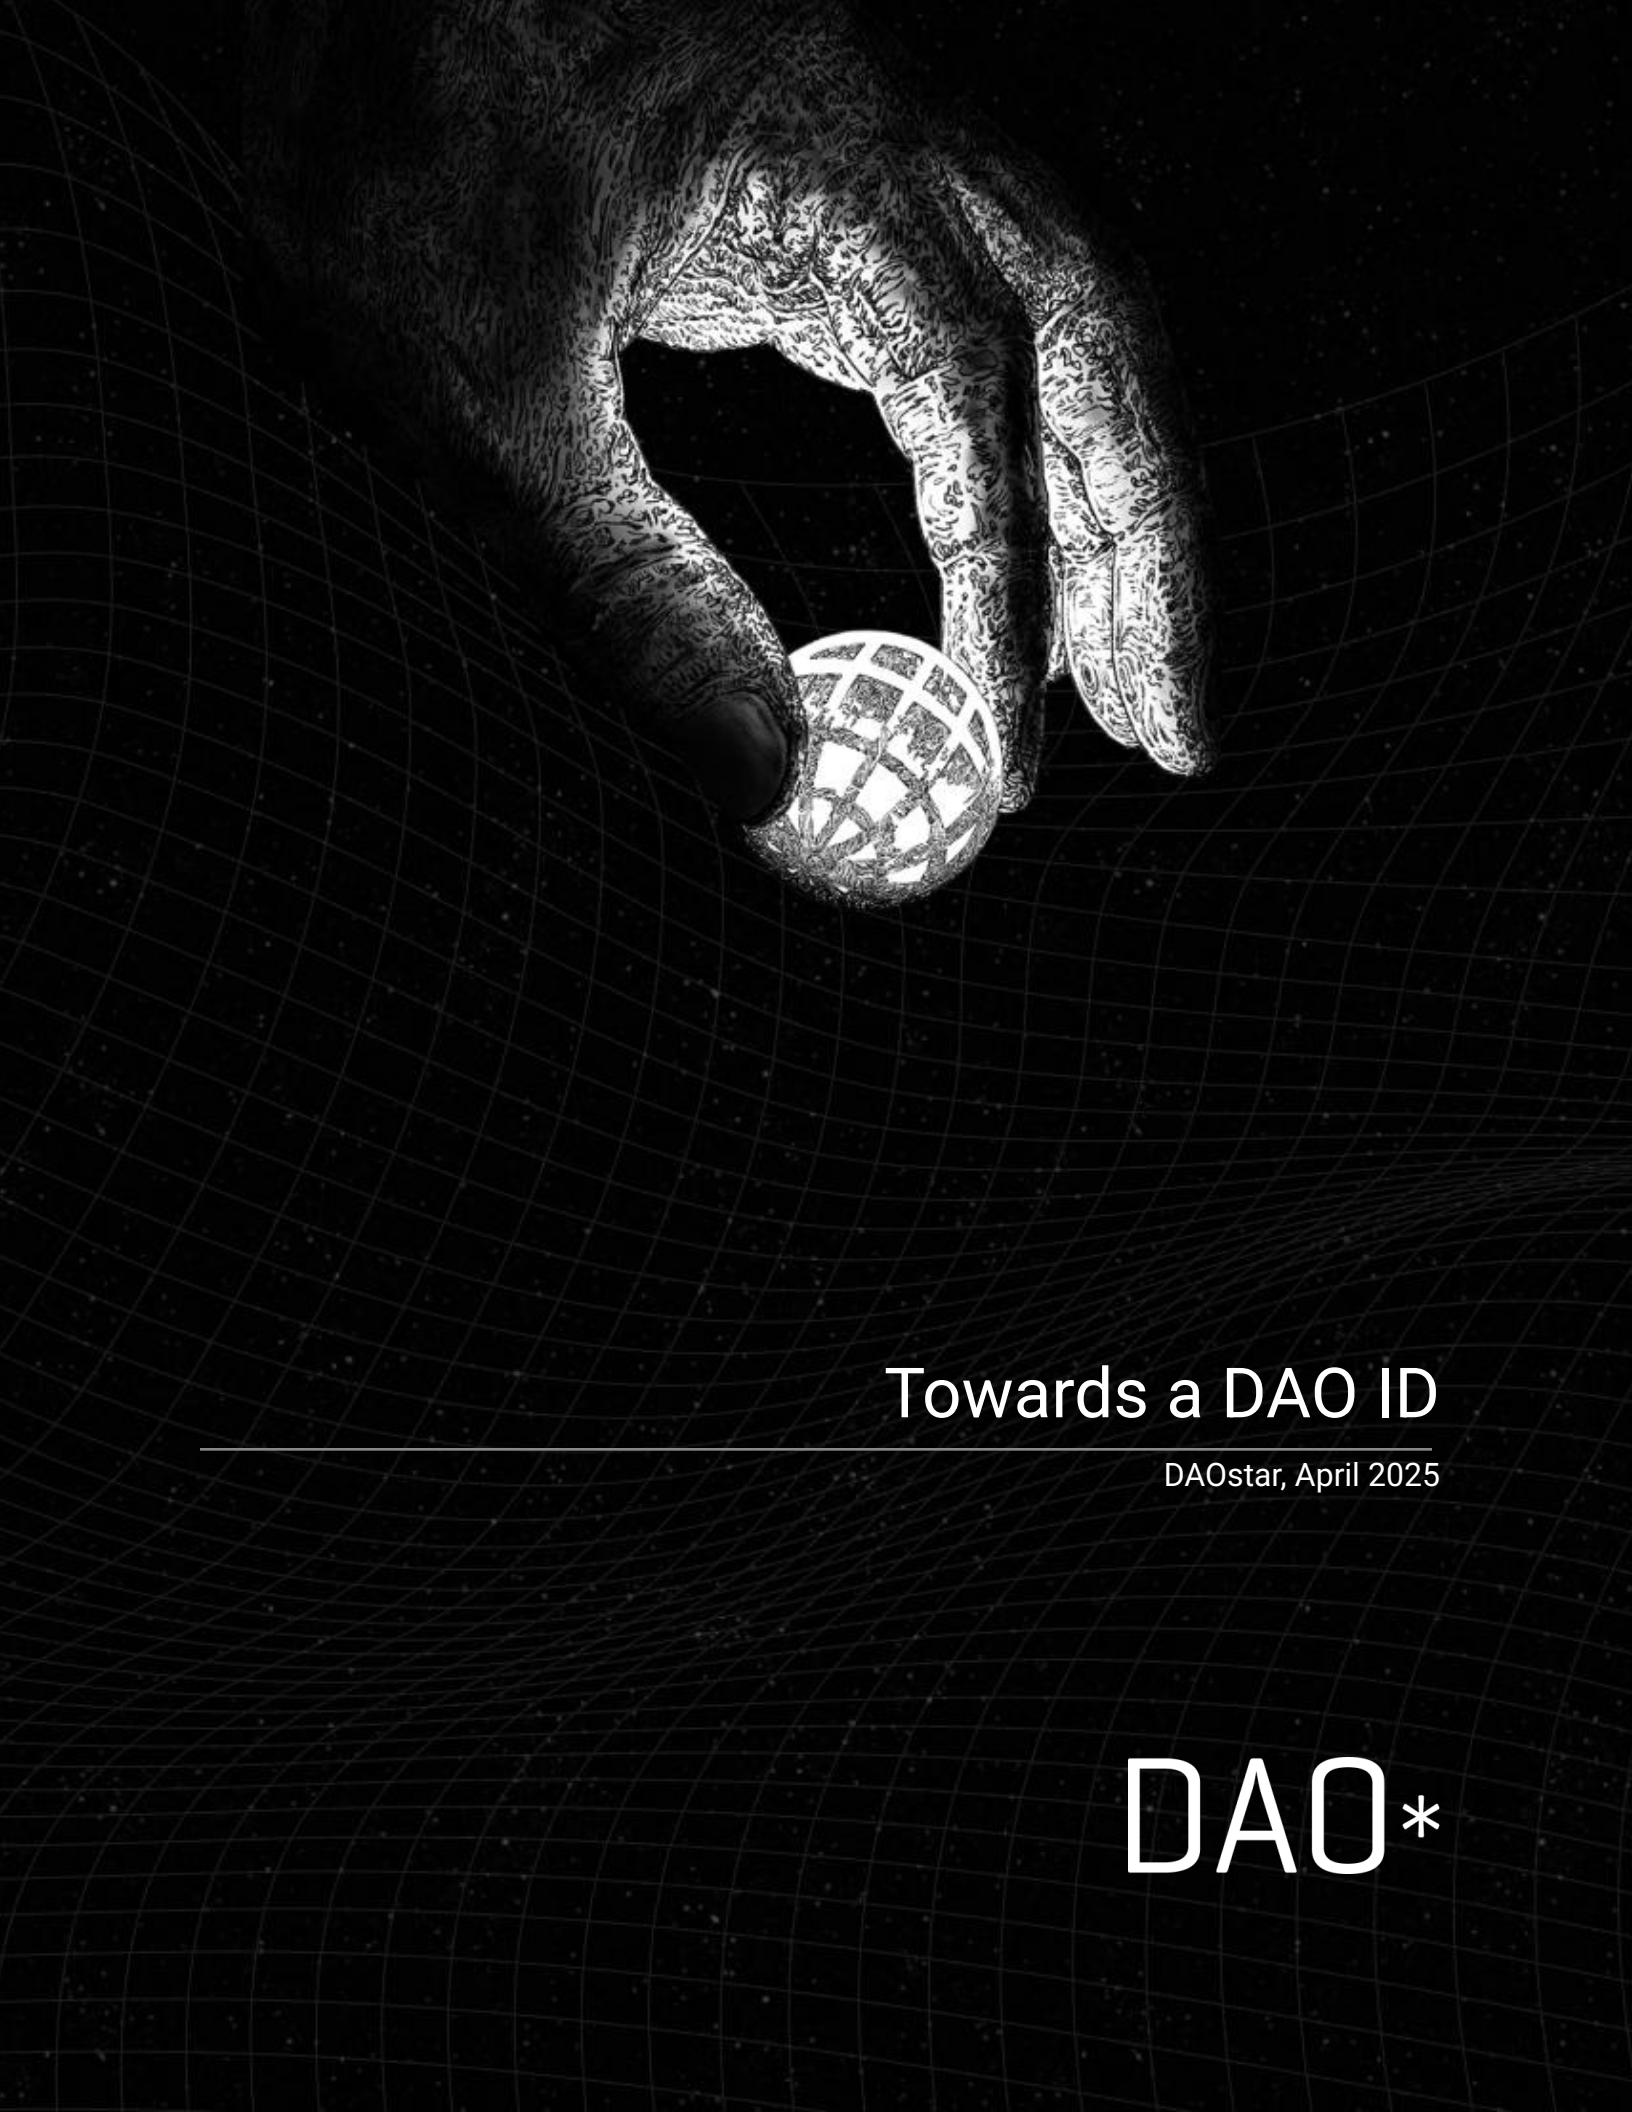

Towards a DAO ID

DAOstar, April 2025

DAO*

Executive Summary

Requirements are explored to extend the current daoURI identity model of DAOstar's identity tooling to a more robust daoID identity model, maximizing for legal recognition and intelligibility even to blockchain-averse authorities. Some limited familiarity with decentralized-identifier prior art is assumed.

About the Authors

Bumblefudge is an independent researcher and consultant working across web3, decentralized identity, and other open-source software research fields. He founded the boutique consultancy learningProof UG.

This report is a publication of DAOstar (or DAO*), the standards body of the DAO ecosystem.

Table of Contents

Chapter I: Problems

After building and widely deploying a Decentralized Autonomous Organization (hereafter “DAO”) toolchain for record-keeping based on the Ethereum standard ERC-4824, DAOstar finds itself a little stumped on how to “v2” its contingent identity system.

After discussing the problem widely with stakeholders inside and outside of that community, Joshua Tan asked me to propose a technological approach to the para-legal identity of DAOs. I do not claim to be an expert in any of the domains touched by this problem set, but at least I’m comfortable proposing solutions to thorny interdisciplinary problems, which seems to be the prime qualification for this kind of work.

I’ll start with my understanding of the main problems that could be solved by a DAO-tailored identifier scheme (be it a DID method, a URI scheme specified precisely enough to be registered with IANA, or some unknown third thing), written in a playful, provocative, maybe even bullying tone. Then I’ll sketch out the shape of a solution and a few high-level approaches, and only then will I turn to a list of candidate solutions. These can be validated and elaborated later by coalitions of the willing-to-prototype that finds this document useful.

Provocation #1: Opt-in Civilization

DAOs are, like blockchains themselves, willful things, contingent, emergent things, rails being laid in front of a moving train by its passengers in real-time. The problem of identifying a DAO is the first step towards defining and stabilizing a subset of all possible DAOs and calling those the “more-recognizable mainstream” of DAOs, whose records stand a reasonable chance of passing muster in most of the legal recognition pathways discovered and expanded in the coming years. Which is to say, a more robust system of verifiable and/or non-repudiable identity statements is a necessary-but-insufficient step towards addressing DAO’s legal trilemma.

Opting into an identification scheme, in the best of cases, would reach a kind of Pareto distribution, allowing 80% of DAOs to sacrifice a little of their freedom and agility to scale up organization, record-keeping, tooling compatibility, and (if I may be so bold) even legal infrastructure appropriate to a novel form of property-first legal person. To optimize for legal recognizability in the design of a DAO identifier system (hereafter “daoID”), I’d like to steelman the most skeptical, anti-DAO critique as the starting point for a design process for the papers that DAOs might hand over to their critics. Trust me, this is the fun part!

> *Upon what basis do persons holding property rights opt into a DAO, and on what basis can a DAO exist to be the referent of an identification to which persons can accountably delegate their property and rights?*

Many people like to start histories of decentralization at the Bitcoin whitepaper, but the savvy critic dodges this framing by pointing out that cryptographic actor models were just a tooling improvement on a **much older category** of legal fictions, namely, those of **“private money”**. By definition, private money wills itself into being, and self-assigns a degree of rights-granting power normally reserved for sovereigns. What matters for our purposes here is that private money goes at least as far back as state-backed money, and is defined in opposition to it. Like a casino on a boat operating under maritime law, private money deliberately occupies a gray zone outside of public courts, and inevitably engenders its own minimalist private (i.e. extrastatal) legal framework.

As the actors granted these non-state-enforced property rights by that self-sovereign money system pile in and start interacting with one another in novel, even experimental ways, the complexity of that system grows and bumps up against its insularity, particularly as dollar amounts snowball and conflicts snowball as well.

> *But how and where to sue your rugpuller, after a private-money deal goes south? How to borrow against an onchain fortune in the offchain world, or vice versa? How to transact onchain about any but the most loosely-regulated of services and intangibles?*

The pioneers of this new world find themselves sorely lacking in legal ground truths. The DAO is a legal person without a court, a stateless corporation, a legal fiction sprung fully formed, Athena-like, from the head of a monetary one. What will be its by-laws, and how can anyone else be expected to enter into contracts or relationships with this mysterious legal Odradek, the legal person composed entirely of unnamed foreigners, which does business under no nation's contract law?

Stated as a design goal, we could say a valuable daoID system must be sufficiently **decoupled from any blockchain** so as to have all its records exported, translated, and legally intelligible *even in countries where blockchains are banned or inaccessible.* It must survive the death, abandonment, or migration away from the blockchain on which it was born to be a viable legal construct.

Provocation #2: Opt-in Globalization

If capital is coded in law, how can global capitalism exist in the absence of a global state and a global legal system? / The solution to this puzzle is surprisingly simple: global capitalism can be sustained, at least in theory, by a single domestic legal system, ***provided that other states recognize and enforce its legal code****. Global capitalism as we know it comes remarkably close to this theoretical possibility: it is built around two domestic legal systems, the laws of England and those of New York State, complemented by a few international treaties, and an extensive* ***network of bilateral trade and investment regimes****, which themselves are centered around a handful of advanced economies.* (Emphasis added. Katherine Pistor, The Code of Capital, A Code for the Globe, Princeton University Press, 2019)

Fundamentally, a DAO is legally an organizational form in much the same way that a cryptocurrency is legally money (and/or commodity and/or security and/or cash-valued coupon): it presents novel forms of documentation, registration, and transparency which are (1) internally cohesive, (2) clearly transact real-world value, and (3) negotiate for some recognition from legal regimes in exchange for some concessions to taxation, regulation, control, accountability, harm-reduction, etc. This "petition" can hopefully scale up over time as more DAOs re-use common mechanisms, data models, and organizational forms; that is, of course, the work of DAOstar and Metagov more generally. But a key assumption in what follows is that these petitions cannot double down on a techno-solutionist appeal to data transparency, particularly if it comes coupled tightly to a specific blockchain (or a set of specific blockchains, or set of protocols used to create blockchains, etc.) as both record-keeping system and source of truth. Such a techno-solutionist appeal would **require the targeted state to legally recognize** that concrete blockchain (or set of blockchains) as an encoding for legal ground truths *before* recognizing the records it encodes and renders verifiable. A DAO has opted in to both *the* blockchain in general, and *a* blockchain in particular, all-but-requiring that government to follow it into both acts of faith for the evidence it brings to the recognition conversation be taken at face value. This kind of "maximalism" is unrealistic and amounts to infrastructural lock-in that any state would rightly be suspicious of, even if it didn't *also* risk displacing the current *Pax Americana* with a *Pax Onchainica*.

A DAO's current records **should not have to wait until** after all the monetary regulatory issues are sorted internationally, and even if they did wait for that baseline of **full recognition**, it would probably not protect the more emergent, "organizing" as opposed to "organized" side of the DAO culture that is of so much value to those trying to resist the total financialization and privatization of all social forms.
Instead, I contend that today's on-chain capital-governance mechanisms need to build (together!) a set of **flexible, robust, and scalable mechanisms that "link out to" and *produce* offchain records that are *already* meaningful and trustworthy** for their specific legal interlocutors, giving them a foundation for concrete proposals that address, e.g., the three

interlocking goals of the DAO policy trilemma. This means tooling and overhead, of course, but it also allows DAOs to earn recognition faster and in parallel to the recognition of blockchains as capital pathways, rather than sitting downstream and waiting for the mountains to move. And the "together" is important, because these mechanisms are expensive to design, build and implement, and because their legitimacy and their chances of lasting fitness-for-purpose both increase the more diverse organizations can design and pilot them together.

On both sides of the national/extra-national divide, much faith (or at least a lot of translation and corroboration) is needed here before such petitions can land, regardless of how legible the outputs of the proposed record-translation is. Not all DAOs will want to be identified in a legally-recognizable way, just as many jurisdictions will not recognize stateless organizations in any form no matter how convincing the analogies. To come back to Pistor's take above on the nature of globalization, most legal systems only recognize the overlap in a Venn diagram of their own legal systems and what they've signed treaties committing them to recognize in the legal systems of friendly nations, and this is often little more than the basic precepts of New York State and London tort and contract law, with a complicated switchboard for negotiating which "local" laws apply in cases of disputes across borders.

A DAO's recognizability, against this geopolitical backdrop, is one of negotiating with states already engaged deeply in **regulatory arbitrage**, and already competing on open-mindedness and acceptance of risk in exchange for investment from mobile stateless capital. These same risk-embracing jurisdictions are also acutely aware of, or even banking on, the limits of that recognition largely ending at their borders, in that only foreign counterparties willing to accept their court as venue for conflict resolution will be doing business with the DAOs they recognize. This means DAO-curious jurisdictions are volunteering to act as a kind of legal "entry point" where a stateless organization can get some tentative or partial recognition qua legal form, which by a lossy transitive property applies to the more conservative, usually richer jurisdictions of the globalized economy via a patchwork of international mutual recognition treaties.

In summary, there is not and will never be (I hope, at least) be a single global standard of "externally legible" targets for translating a DAO's records to, even if New York State and London property and securities law come worryingly close to being that in the present global order. Instead, the translation of a DAO's records, membership, holdings, governance, etc., will always be a dynamic and open-ended patchwork. An "identity" made up of non-repudiable, public claims (and probably even some provable negative claims, i.e. commitments that exclude private counterevidence!) is, indisputably, necessary but insufficient ground upon which to make sure claims to existence and exertions of rights fit through the moving goalposts of a global competition for regulatory arbitrage we could call the international "fine print" Olympics.

Provocation #3: Identity Within and Beyond a Blockchain

Using a blockchain to self-manage identities is not enough to make an identity fully "self-sovereign" (Src: Me, 2018)

There are many ways in which a blockchain is a less unitary and definitive ledger of social consensus than people like to admit. They are great for providing ground-truth about the movement of assets over time, but terrible at providing ground-truth about the actors in the system, since most of them enforce almost nothing about the binding between asset-holding actors and real-world/off-chain persons. Each layer on which a blockchain identity is contingent adds to the cross-chain and off-chain contingencies piled up behind the word "identity". I'll walk through these layered contingenies quickly as a kind of to-do list, which any reasonable blockchain-based organization would have to address when making itself externally legible to a blockchain-critical world.

The most obvious of these are so-called hard forks, whereby the community of actors and participants in a cryptographic consensus system splits by a kind of cellular mitosis and produces two competing version of the same blockchain (in more technical terms, two mutually-invalidating branches of one monotonic linkedlist of shared state). The early years of blockchains as economic systems saw them compete in relatively primitive and direct ways, interoperating as little as possible and spawning forks off of "mainline" Bitcoin one after the other every time a minority reached critical mass to survive a secession. An identity rooted in a blockchain can only be as stable or unambiguous as the identity of the blockchain itself: furthermore, there needs to be a strict 1:1 mapping of blockchain identifier to the ledger or cryptographic substrate vouchsafing it, at least for any bounded period of time whose records will later be credible.

When the Ethereum community reached social consensus to **intervene directly with the chain-state** (and thus the "source of truth") to reverse the damage of the famous "DAO Hack," this majority decision caused a particularly contentious and ideological hard fork in early Ethereum's fledgling community. If Bitcoin's many forks and carbon copies showed the fragility of userbases and social contracts, the genesis of Ethereum Classic shows that even within the opt-in construct of a blockchain, coalitions of the powerful can move mountains and reverse commitments. What the DAO Hack remediation achieved by extra-technical means, a garden-variety 51% attack can also achieve within code: the data is only as good as the consensus around its production.

There is an obvious, commonsensical appeal to registering an organization on the global state machines powering 24/7, always-on capital markets. But some of those global all-powerful state machines (*cough, cough* SOLANA) just **go down sometimes, for a day at a time**. Others have their node-to-node networking intercepted by Great Firewalls or cyberwarfare brown-outs, and the oceans haven't even really started boiling yet. How does an on-chain entity square its

internal records with what happens off-chain during, say, a daylong blackout, or a continent-wide brownout?

In the traditional world of finance, locking up assets as collateral to borrow something else, then locking *that* up as collateral, or speculating with cash borrowed against unpaid but legally-binding invoices, or other forms of risk-multiplying leverage stacking, is called **hypothecation** and it is tightly regulated and monitored to contain its worst case scenarios from creating domino effects in the capital markets. On blockchains it's just called "DeFi" and glorified as a value unto itself. Those domino effects can cause prices to swing so much that even banal things like "gas" (per-transaction usage costs) and transaction speeds are affected for unrelated co-users of the same system (like DAOs or DIDs).

In this sense, what would be legally useful to get asset-handling DAOs recognized legally is a kind of **cross-chain identity** and **clearly-defined rules** for how to balance assets across multiple chains, which double as a public commitment to manage risk with automatic rebalancing mechanisms across venue-diversified treasuries. This would make treasury/asset-checks and solvency audits much more complex to execute, but also more recognizable to legal interlocutors as sufficiently distanced from any one blockchain's (or DeFi protocol's) inherent risks. Making such checks aware enough of lock-ups, staking, bridge-locks, time-locks, etc etc is a tall order, and hopefully orthogonal here, but I mention cross-chain assets because a necessary (and insufficient) condition for such checks is an identity system that declares (hopefully verifiably and non-repudiably) which assets are controlled on multiple chains over time, and in multiple kinds of locking contracts such that any legal situation requiring entities to prove their solvency or expose assets for audit can check such assets against public records rather than just-in-time disclosures.

Maybe there is no other way to incorporate all of this generically enough in an identity system than to add one more optional array of URIs, but it's worth mentioning if we're listing MUSTs and SHOULDs. Maybe many of these failure-modes are already accounted for in the legal systems being addressed, but keeping them in mind when designing mitigations and translations ensures there is some foothold for triaging these issues as piecemeal recognition proceeds and travels to complex global corner-cases.

Provocation #4: Ships of Theseus All the way Down

Even vanilla legal entities are leaky abstractions over leaky ships of Theseus.

Delegation (i.e. who can actually represent a collective at any given point in time) is as tricky a technology to specify in contracts and statutes as it is in data management systems. In fact, the latter inherits most of its mental models and terminology from the former, as anyone deep in the literature on authorization systems can tell you. If technological delegation systems seems rigidly and needlessly hierarchical, it's largely due to mental models inherited from law which seeks to minimize discretion and ambiguity. When parties start disagreeing on who could represent what collective when, the fallback and ground truth is to look up public documents and (recognized, verifiable) public claims to trace who has the most leg to stand on in the dispute; crucially, this includes named officers of incorporated entities and board members of public entities. In incorporated entities, this is easy enough– named officers and board memberships are pretty clearly public and available information, and any time membership changes need to be published, the burden of publication is on the currently-public members. Reaching **parity with formal, incorporated entities** requires similar (but self-driven!) formalization.

While it is debatably a nice-to-have rather than a MUST, we might want the URIs naming "officer"-equivalents and liable members to be checked not just against current state on one blockchain, but also historically against multiple oracles and even, perhaps, certain authoritative sources offchain. This is simple enough to do when all canonical records live on the same [stable, never-down, long-lived, fully-public] blockchain, but any multichain or off-chain mechanism would also need to be just as historically verifiable, at least to the level of granularity of most-recent block at any given timestamp.

Furthermore, even setting aside multi-chain gaps in logic, delegation systems too tightly coupled to *any* on-chain logic (native or custom) tends to be less legally intelligible because it requires technical expertise and historical research to translate to legal fact, based on what exact version of client software was running network-wide at the time of that block, etc etc. Auditable, standards-based delegation languages produce more unambiguous audit logs, which are cheaper and easier to translate into evidence; the next best thing is chain-aware groupware that translates decisions and delegations to a neutral language at the time, in case of disputes years later when chain-specific logic becomes an expensive historical matter to resolve.

Ironically, delegation has long been the biggest design challenge in finding the translatability and equivalence *between* DID methods. It has also been a kind of overwhelming quagmire in the design of some DID methods *not* locked to a single blockchain, where the field is wide open and more choices can be something of a curse to designers and communities, choosing between how to accommodate multiple of the world's online and offline delegation languages and authorization event-logging systems. Which brings us to the most vexing provocation of all:

Provocation: What does anyone even mean when they say "DID"?

Depending on whom you ask, a DID method can be:

1. a **URI scheme with a resolution mechanism** stapled to it
2. an **export format** for identity documents allowing actors outside that that system to deterministically and transparently fetch the current document for a given URI
3. a **bag of keys** and, optionally, endpoints (both serialized as URIs) that can be fetched for any specimen of a given type of actor in a DPKI (decentralized public key infrastructure) system
4. a specification formalizing how to **dereference a URI to**… a document containing **more URIs** (i.e. a slightly more kitchen-sink URI scheme)
5. an abstraction allowing a **cryptographic identity to mutate** (i.e. rotate, grow, and shrink) verifiably over time
6. a user's manual for a **censorship-proof and resilient** identity-document substrate
7. the constitution for a service that dereferences URIs to **certified documents**, whether by delivering self-certifying documents or by certifying the documents it delivers

While this might sound like hyperbole, I have attended in-person meetings of the DID Working Group over the years, and each of the above caricatures is based on the working definition various key members of that working group hold and act on (and, debatably, pull the rest of the working group towards over time). There *has* been some progress over time to tighten this loose bag of purposes, such as my former boss Wayne Chang's ideological intervention arguing for a feature-matrix approach to DID interoperability and migration thinking.

But fundamentally, we're still here, trying to figure out what generalization is worth making about DIDs, when there is such diversity between their URI schemes and their resolution mechanisms. **Each DID method is grounded in a radically different *kind* of ground-truth-producing system**. They're not much more generalizable or schematized than DAOs– if anything, they are even more emergent on both the technological and social/business layers. There is no general sense in which "DIDs" provide a path to meeting all the requirements above, nor are DIDs easily classifiable into a few categories, which each drop one or two requirements from the list to deliver the rest. That said, the "trait" methodology mentioned above that some are using to differentiate DIDs by feature-sets might still be useful in finding 1 or more DID methods that are useful enough to be better than status quo, and for writing a specification that names everything else a better DID method would have. That's what I have done, at a high-level, below, and could do more systematically after requirements have been clarified and validated by a set of DAOs interested in building common identity and administrative tooling together.

Historical Verifiability of Mutable [DID] Documents

Historical resolution, long a contentious and optional feature of DID methods, is actually only supported by a few of them. Most DID method designers punt this extremely difficult set of corner cases to distinct “internal” or orthogonal mechanisms (i.e., independent of the DID document’s resolution), rather than make available to any consumer of the DID document detailed information about what a given DID document was at a specific point in time via the optional `versionTime=` query parameter defined in the universal DID URL API. Indeed, even query parameters themselves are kind of rare in DID-based systems, which mostly do not further specify or require usage of DID URLs, so limiting ourselves to DID methods that handle historical verifiability automatically is a fairly constraining decision without many options, and should be compared (apples-to-apples) with solutions where such verifiability is guaranteed on another layer (e.g. historical chain-state).

Chapter II: The Shape of a Solution

To decouple a bit the question of whether or not the problem space is best addressed by one or more current or future DID methods from a higher-level analysis of the space itself and the requirements, I have elaborated some families of options first, relative to daoURI status quo. In doing so, I made a solid assumption that a daoID defined a DAO over time without taking authority or priority over the daoURI valid at any point.

Thinking through the provocations above to generate a list of functional requirements and properties that a daoID system would need to maximally scaffold legal recognition from a skeptical and offchain legal authority, we could enumerate these requirements approximately as:

Summary of Requirements above

1. Opt-in upgrade: no current ERC-4824 DAO should *need* to implement or change anything if the kinds of legal recognition or non-repudiable record-keeping enabled by this proposal are not urgent for them.
2. Some kind of historical verifiability of previous versions of the daoURI document.
3. Either a long-term identifier per DAO OR an “oracle” that can return all previously-authoritative daoURIs over time when queried with any current or historical daoURI (at least for all DAOs opting into such a comprehensive oracle, as per req. #1).
4. A chain-agnostic/off-chain daoID Document, if established, needs to be verifiably and bi-directionally linked to all the relevant on-chain daoURIs and off-chain documents verifiably linked from on-chain (directly or indirectly).
5. Extensible enough to verifiably link to additional novel Document types, so as to accommodate additional reporting, auditing, authentication, deanonymizing, and other mechanisms as needed for legal recognition (particularly ones where underlying blockchain records and identities are categorically inadequate and external notarization or machine-readable equivalence statements are needed)
6. Unambiguously resolvable in case of blockchain forks or outages or other availability issues
    - (May require explicit disambiguation logic around forks as force majeure)
7. Clear fallback for DAO-Hack style intervention in chain-state
    - (May require explicit deferral to or from altered or rolled-back chain-state)
8. Multichain and Multi-VM dereferencing: child-DAOs, members, and controlled accounts might all live on other chains than the daoURI, and these must all be expressable at least in the form of a daoID, if not in the legacy daoURI itself
    - (May require significant work to support non-EVM chains over time)

Relationship to the daoURI() interface

Realistically, a daoID is only worth offering as an extension of ERC-4824 and daoURIs as they exist today. Stated technically, we could say that a worthwhile (and realistic to deploy) daoID would need to be a strictly backwards-compatible superset of daoURI, i.e., every daoURI returned by the daoURI() on-chain calls possible today MUST be a valid daoID value. Practically, this means that a daoID is just a BIGGER bag of URIs, which might add some optional properties if there is value in adding them, resolved in a different way than just querying the same chain with a different call. That passive-voice “resolved” does a lot of work, though! DID methods, in the best of cases, abstract over [pre-existing] resolution methods to bridge consumers and producers inside and outside of them, whereas here any resolution method chosen is a novel infrastructural commitment demanded of every DAO self-publishing these records in addition to their authoritative daoURI records on their [primary, current] chain.

Whether the solution involves DIDs or not, the publication and resolution mechanisms for these chain-agnostic documents make or break the feasibility of any such proposal.

daoID Document as Microledger of daoURI Documents

Working backwards from failure modes, one backstop that strikes me as a hard requirement to commit to upfront is for a given daoID to be considered valid IF AND ONLY IF it also is or contains the current daoURI document. Extending the daoURI document system is already a huge ask of daoURI users, without also creating a risk center and attack vector in case the extension gets out of sync somehow with the simpler, easier-to-maintain legacy form.

This property could be achieved two very different ways: making a backwards-compatible version of the daoURI method that returns additional properties and all the current ones, OR making a daoID that lives somewhere else, includes a link to the daoURI, and is not considered complete until BOTH documents have been dereferenced (and perhaps merged). I refer to these two variants as “second-document” approaches, above and below.

From simplest to most complex resolution methods, I can think of three families of solutions.

Solution-Shape A: Back-Linking daoURIs

A single property could be added to each daoURI document called, say, `previousVersion` which pointed to the most recent previous state that the daoURI resolved to (if addressable as a standard URI, or at least to the transaction that updated it, if this can deterministically produce the previous version). This would mean that from today's daoURI, one could walk back across all updates to inception and have the history of each Dao as a "microledger" of daoURIs. Since each finite update to daoURI doc and/or memberships or other important properties occurred in a timestamped (or at least block-numbered) transaction, fetching historical state for a given daoURI at a point in time would be deterministic, if slightly networking/compute intensive without a comprehensive indexer, historical node, or at least a pre-loaded cache, of entire histories of each daoURI.

Note: this would only really be feature-complete for net-new DAOs, since the microledger would end abruptly at the first post-implementation update for any DAO whose history stretches before the implementation.

Unlike the two alternatives below, where the legacy document and resolution is extended by a second document type and resolution mechanisms, requirement 4 is moot here because no second Document-type is added. By iterating the core daoURI mechanism, though, all changes to support the other requirements would have to be made together, along with the backlinking change, in one atomic upgrade. For requirement 5, this would simply consist of additional properties, but solving for 6, 7, and potentially even 8 might not be possible in this solution-shape, or at least require validation research in a different direction.

Solution-Shape B: Off-Linking from daoURIs Documents to daoID Documents

Similarly, instead of `previousVersion`, a property like `alternateDocument` could contain 1 or more URLs pointing to an “extended” version of the same document, which also includes metadata about update history, links to previous or individual versions, etc. This could be content-identified, i.e. `https://{ipfs-cid}.subdomain-gateway.ipfs.io` or `ipfs:///{ipfs-cid}`, or `at://{did}/{lexicon}/{cid}`, or some other content-addressed, high-availability URL, in addition to or instead of conventional HTTPS urls.

Since this off-chain copy of the same document would probably be less constrained for space than something on-chain, it could conceivably take a different (i.e. more complex, even layered) structure than merely adding a single backlink. For example, the daoID could refer to an intermediate document containing links to every version of the file sorted in reverse chronological order, tuples of links and metadata, triples of HTTPS link, alternate link, and metadata, etc etc. Whether or not this added complexity here makes sense depends on the trust model and on the degree to which it solves the legacy problem of daoURI document historical guarantees pre-upgrade.

Mitigations for requirements 6 and 7 can be encoded quite simply in parsing rules for off-chain documents as a “state of exception” to requirements 1 and 4, i.e., “until outage/fork/consensus issues/availability issues are resolved to X standard, most recent on-chain daoURI Document remains valid and off-chain extension document can be acted on.”

Solution-Shape C: Completely Separate Resolution Mechanism

If, instead of changing anything about the daoURI Documents, a second step were introduced *after* resolving the daoURI, no change would need to be made to how daoURIs are published and resolved. For example, with the current daoURI document in hand, one could take the daoURI (or a hash of the current daoURI document, or any other deterministic transformation of it into a URN) and resolved that in some alternate registry (onchain or off-), to get back the "extended version" of the Document, or all other documents pertinent to, linked from, or linking to that Document.

"Deterministic" is doing a lot of work in the sentence above, because the universal determinism of [canonicalized] daoURI documents to URNs is required throughout the system for them to be reliably used for discovery, deduplication, and subsequent authentication (i.e. to function as a checksum). Note that a classic "just use IPFS" approach is probably insufficiently deterministic here; for content-identifiers to be load-bearing and error-proofed in this way, something more specific (like serialization-specific canonicalization of documents and generation of CIDs according to a precise IPFS profile)

Mitigations for requirements 6 and 7 can be encoded just as simply as above in the parsing rules for off-chain documents as a "state of exception" to requirements 1 and 4.

Even if append-only, even if content-addressed, this approach also brings up the usual governance questions about how a common corpus or registry could be governed in common, since it would be a public good and not free to maintain, who can be incentivized to persist "fallback copies" à la Internet Archive, etc etc.

Conclusion: Future Work

In the document above, I have striven to describe neutrally both sides of every tradeoff as a family tree of options. Realistically, builders of tooling want actionable starting points for prototyping, and as the DID space is crowded with prototypes and short on authorities and marketing budgets, very few observers can even point to a list of relevant prior art.

My recommendation for researchers and prototypers is to think through the following five approximate architectures as the five ways of making a DID-based (or DID-like) that could be evaluated against the preceding requirements and criteria.

- Candidate Solution: Minimalist Upgrade to daoURI (only add backlinks and minimal/open-ended extensibility)
- Candidate Solution: ENS-based second-document
- Candidate Solution: A Ceramic-based broadcast/self-publishing network
- Candidate Solution: A `did:dht`, pkarr, or `at://` broadcast/self-publishing network
- Candidate Solution: A `did:webvh` (i.e., HTTPS-based) microserver per DAO, tracked by multiple independent mirrors/“archival nodes”

Acknowledgements

Thanks to the Decentralized Identity Foundation for fostering much of the prior art referenced here, and of course to the DAOstar community for justifying the effort.

Bibliography

1. Ethereum Improvement Proposals. (n.d.). *ERC-4824: Common Interfaces for DAOs*. EIP Tools. https://eip.tools/eip/4824
2. Metagov. (2024). *The DAO Policy Trilemma* [PDF]. GitHub. https://github.com/metagov/daostar/blob/main/Reports/The%20DAO%20Policy%20Trilemma%20(April%202024)%20v0.1.pdf
3. Pistor, K. (2019). *The code of capital: How the law creates wealth and inequality*. Princeton University Press.
4. learningProof UG. (2018, October 28). *Self-sovereignty and autonomy*. https://learningproof.xyz/self-sovereignty-and-autonomy/#sovereignty
5. Graeber, D., & Wengrow, D. (2021). *The dawn of everything: A new history of humanity*. Farrar, Straus and Giroux.
6. Lopatto, E. (2022, December 9). *Who's snitching on the big crypto group chat?* The Verge. https://www.theverge.com/2022/12/9/23502193/ftx-alameda-binance-kraken-tether-exchange-texts
7. Chang, W. (2022, July 1). *Upgradeable decentralized identity - DID method traits*. SpruceID. https://blog.spruceid.com/upgradeable-decentralized-identity/
8. Decentralized Identity Foundation. (n.d.). *DID traits*. https://identity.foundation/did-traits/
9. ricmoo.eth. (2017, May 17). *ENSIP-5: Text records*. ENS Docs. https://docs.ens.domains/ensip/5
10. Decentralized Identity Foundation. (n.d.). *did:dht method and server implementation* [Source code]. GitHub. https://github.com/decentralized-identity/did-dht
11. pubky. (2025). *pkarr: Public key addressable resource records* [Source code]. GitHub. https://github.com/Pubky/pkarr
12. BlueSky. (2023). *AT URI scheme specification*. https://atproto.com/specs/at-uri-scheme
13. Decentralized Identity Foundation. (n.d.). *DID:WebVH*. https://identity.foundation/didwebvh/next/