Taxonomy guide: noslegal v 3.0

31 March 2025

This guide explains:

1. Who the noslegal community are
2. Who our taxonomy is for
3. What you can use it for
4. Our approach to taxonomy design
5. Implementing the taxonomy
6. History, governance and roadmap

For the taxonomy content, please refer to the separate **release notes**.

1. About the noslegal community

We're a voluntary community of lawyers, engineers, legal business and knowledge professionals, and more. Although we're based mostly in the UK, we have an international perspective. We're currently actively introducing noslegal to continental Europe, and keen for involvement from elsewhere in the world.

We produce data standards and guidance for use in law and legal work. So far, most of what we've done has been focused around a taxonomy.

We emphasise the principles below in all our work.[1] For our taxonomy work specifically, some further design principles are outlined in section 4.

The five principles

- **Open source.** Anyone can use, study, modify and redistribute the things we publish, for any purpose and without charge.[2]

- **Community-based.** Anyone can contribute to the things we publish, though we have governance to ensure quality. We also emphasise community as an end in itself.

[1] In addition to taxonomy, we anticipate publishing in the near future some schemas (i.e. the attributes which documents, matters and people may have, drawing on the noslegal taxonomy among other things). We also expect to publish more guides and articles on various topics beyond taxonomy in future.

[2] The specific licence we publish our taxonomy under is Apache 2.0 - a widely-used permissive open source licence mostly used for software but which can be used more generally for works of authorship such as the noslegal taxonomy. The taxonomy is published at https://github.com/noslegal/taxonomy

- **Useful.** We do things of practical benefit, not just theoretical interest.
- **Simple and modular.** We produce material in simple, digestible modules. But these can be combined with each other, remixed and extended for flexibility and to capture detail.
- **Broad.** Our work should be relevant to a wide range of organisations and needs. Although national or local initiatives are possible, our approach overall seeks to be internationally relevant rather than based on a single legal culture.

2. Who the noslegal taxonomy is for

2.1 Three types of organisation

While the law is typically created, applied and enforced by authorities, specialised intermediaries strongly influence how it operates in reality, including:

- **Legal services providers.** Lawyers and others who practise law and provide related services. This includes traditional professional organisations such as law firms and barristers' chambers, as well as not-for-profit providers such as law centres. In recent decades it has also come to include corporate enterprises known as alternative legal services providers (ALSPs).
- **In-house legal services.** Businesses and larger public bodies typically have legal departments. Typically these are staffed mainly by lawyers, do some legal work directly and manage the sourcing of work from external legal services providers.
- **Legal technology companies.** These provide software and, in some cases, information used either by the first two intermediary groups, or directly by end users (such as individuals, small businesses, or others who cannot afford legal services providers).

The noslegal taxonomy is designed to support these three types of organisation.

2.2 Four areas of need

Legal services providers and in-house legal services have four major overlapping areas of need to which the taxonomy is relevant:

- **Opportunities.** Determining and influencing what work is upcoming or available. Within a legal services provider, it includes marketing and sales. Within a legal department, it includes developing internal relationships

with the business, identifying their needs and what's coming up; determining how to balance different types of resourcing, then arranging the sourcing of that work.

- **Delivery.** This involves scoping, planning, pricing, executing and managing the work, including subsequent adjustment of plans and pricing if need be.

- **Knowledge.** This involves making useful information available to help with delivery, which may draw on the organisation's collective experience, or external resources.

- **People.** These needs include determining and communicating what experience and skills an organisation's people have, and determining what's missing so that gaps can be filled by recruitment, training, secondments, work allocation, contracting-in or in other ways.

These four areas overlap and interact in ways that make it desirable for the higher-level concepts used in each to correspond.

For example, a law firm will ideally target its marketing and sales activities (**Pipeline / Sourcing**) on areas which can generate significant amounts of worthwhile work (**Delivery**) and will make relevant learnings as to how to do so effectively and profitably (**Knowledge**) and generate credentials (**People**) which can support marketing and sales (**Pipeline / Sourcing** again) and help those doing the work in future (**Delivery** again and **People**).

Organisational reality is that the core functions in these four areas are often handled by separate groups within many legal services providers, each with its own characteristic software applications.

- **Opportunities:** in legal services providers, typically handled by lawyers together with - in larger firms - specialists in marketing and sales (widely known as "business development" in such firms) and sometimes other specialists (e.g. pricing). Customer relationship management software may be used, often in law firms with a particular emphasis on capturing credentials. In legal departments, typically handled by lawyers, sometimes with the involvement of procurement functions - though for legal work, these are often less involved than in other corporate purchases. Vendor management software may be used.

- **Delivery:** generally handled by lawyers and other "fee earners", sometimes (in larger organisations) with specialist input in areas such as

pricing and project management. A lot of this work is still done in email and documents, but there are various types of software for managing it in a more systematic way. Finance systems used in law have specialised features, in particular to deal with the often heavy reliance on time-based billing.

- **Knowledge:** often handled by specialist knowledge management / information services or, in smaller organisations, a library. There may be specialist software for capturing and making available forms of knowledge internally, though external services are very significant for meeting knowledge needs.

- **People:** mostly handled by human resources with specialist software (or at least spreadsheets) but also with major cross-over with marketing (e.g. website credentials) and sales (e.g. credentials for pitches) and knowledge management (e.g. some types of training, though HR may take the lead on other types).

In increasingly complex and competitive markets, there is growing recognition that it is desirable for data generated in each area to be understandable and reusable in the others.

For example, doing legal work generates financial data, credentials and knowledge, all of which can be reused to win more work, and do it to a higher standard and on a more attractive financial basis in future. Similarly, buying legal work can, if managed appropriately, generate data on the quality and value for money with which the work is done, as well as knowledge, all which can help with sourcing and management of future work, and in avoiding unnecessary work.

Traditionally this has not been done systematically. Working to a shared high-level taxonomy is an important aspect of being able to achieve this, and is more manageable than a constant mapping and remapping exercise between independently managed taxonomies.

Needs will diverge at lower levels of detail. For example, a knowledge management team will typically have more need for more granular legal concepts and sub-concepts than a marketing team. This can be handled by extending the taxonomy to meet such specialist needs.

2.3 Who the noslegal taxonomy is not designed for

A significant limitation of legal taxonomies is when they become overfitted to a particular place or needs. As the noslegal taxonomy is designed to help address practical needs in an international context, we have emphasised throughout that it is not specifically designed for:

- Legal authorities such as courts, regulators and law enforcement.
- Legal research.
- Teaching.
- Theoretical or academic law.

This is because these areas tend to make jurisdiction-specific and legal-conceptual demands which would undermine the value of the noslegal taxonomy in the areas for which it is designed.

This constraint is not intended to be negative or dismissive. People working in these important areas are free to use, modify and extend noslegal in their contexts. But just bear in mind that, unlike many legal taxonomies, noslegal has not been designed with these needs specifically in mind.

3. What the noslegal taxonomy can be used for

3.1 Search and analysis

In broad terms, the noslegal taxonomy is useful for search and analysis.

- **Search** means finding specific things.
- **Analysis** means understanding the characteristics of groups of entities (such as matters, documents, organisational groups and individuals) - whether numerical (e.g. financials, numbers of documents) or qualitative (e.g. the experience someone has in a particular type of work) or some combination (e.g. what % of their recorded time someone has been spending on a particular type of work).

Examples of search and analysis uses:

Area	Search	Analysis
Opportunities	Have we done any matters like this recently that we can mention as credentials in this pitch?	What conversion rates are we seeing for different types of work and what do we need to improve?
Delivery	How is fixed pricing likely to work out on a matter like this - what has happened previously?	What kinds of work have we done for this client this year, in quantitative terms? What types of work are most and least profitable? What kind of work suffers most from budgetary overruns and write-offs, so we can manage those more effectively in future?
Knowledge	How can I do X (of immediate importance to me in this matter)?	How proportionate and up to date are our knowledge resources compared with their importance to our profitability?
People	Do any associates have capacity and experience in this kind of matter ?	Are we spreading experience out widely enough to support people's development and the firm's effective handling of work?

While addressing these topics individually is possible, it can be done much more powerfully if the data is understood within a common taxonomy. And, crucially, the taxonomy need not be particularly detailed for that to be useful. What is certain is that a relatively simple taxonomy applied accurately to data supports these topics better than a more detailed taxonomy applied inaccurately.

3.2 The noslegal taxonomy and AI

In addition to the search and analysis uses discussed, we are also seeing increasing recognition within the noslegal community and beyond of the relevance of strong taxonomy and data classification to support AI use cases.

Our view is that:

- A strong human-designed taxonomy such as this is extremely useful if using AI, as it defines what is valuable to identify and distinguish.

- You can certainly use AI to help classify things, to use the classified things, and to travel the semantic "last mile" not covered by the taxonomy.

- It's also possible to use AI to help develop the taxonomy. We have used AI tools to ideate and critique aspects of the 2025 release. We expect in future to explore this further, together with the related topic of mapping noslegal to other taxonomies where relevant to do so.

3.3 Standard or starting point?

Our taxonomy can be used in a few ways:

(a) Simply adopt all or large parts without modification. This is straightforward in the sense that the different parts have been designed to work as a whole. This approach also makes it easier to accept noslegal updates in future. And if adoption of noslegal continues to grow, you can start to address sharing data more effectively and efficiently e.g. between clients and law firms.

(b) Start with the noslegal standard but extend it (rather than modify it) with further levels of detail or further facets to meet your specialist needs. This can be a good way of achieving the benefits of (a) without the constraints.

(c) Adopt or continue with a custom taxonomy, but draw on noslegal for ideas or concepts to patch or improve your taxonomy, or even to show internal stakeholders the possibilities. Many of us involved in noslegal have worked with taxonomies in large, complex organisations, including recent, current and planned implementations. Suffice to say that it is sometimes unavoidable, but the complexity of doing this effectively and managing it over the years, across functions and locations, can become challenging even if you have a specialist team to help do so (and if you don't have such a team then the risks are magnified).

(d) Start with noslegal but then substantially modify it (as well as extending it). Depending how far you take this, it may in substance be more like (a)/(b) or more like (c).

Whichever approach you take, we would suggest approaching the topic consciously, considering benefits and risks over years, not just the short term.

4. Our approach to the taxonomy design

We mentioned in section 1 above the 5 principles which we follow. This section describes our approach to the taxonomy design in more detail.

4.1 What is a taxonomy anyway?

A taxonomy is a 'tree' of concepts addressing a certain topic:

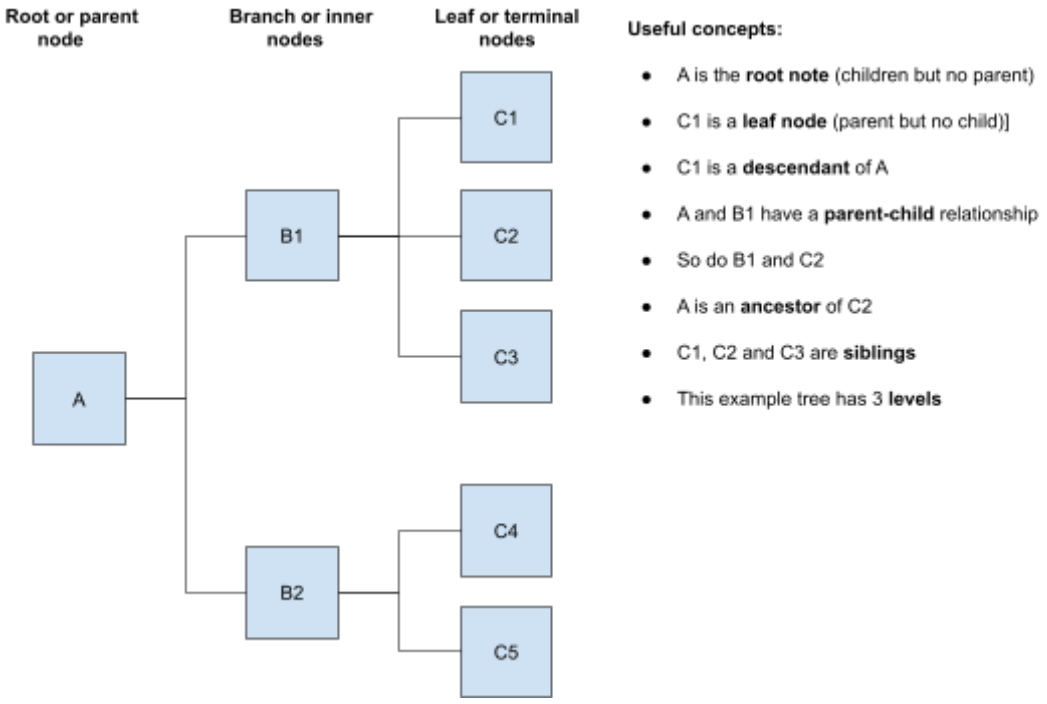

Diagram 1

We concluded early on that a taxonomy is the 'sweet spot' for the practical purposes discussed in this guide, compared with:

- A simple list of words, definitions, synonyms and examples (i.e. a dictionary or "controlled vocabulary").

- An ontology[3] — which can capture a broader range of relationships but which requires a greater strategic commitment and investment than law-focused organisations have typically been prepared to make. Their existing specialist software systems are at present not built to handle ontologies and knowledge graphs, and even large firms have thus far found attempts to introduce them rather daunting. A taxonomy allows for

[3] This is an information science-specific usage of the word "ontology" explained here — not to be confused with the philosophical usage.

progress within the law firm and legal department status quo, without allowing 'ideal' to be the enemy of 'better'.

4.2 Monohierarchical

The noslegal taxonomy is monohierarchical: this means that each sub-concept can only have one direct parent concept.

When combined with faceting (below) a monohierarchy allows for a sophisticated description of relevant 'things' while being simpler to understand and implement. That said, we do in some parts offer:

- alternative aggregations in Places (EMEA / APAC / AMER) which could be implemented (in addition to Europe, Asia etc) as a polyhierarchy. But you could also just choose one rather than both; and

- some relationships which are many-to-many, e.g. a single country may be a member of both the New York Arbitration Convention and the Hague Child Abduction Convention. But in these cases we have used flat lists (e.g. Conventions) rather than taxonomies on one side of those relationships.

Why not a polyhierarchy?

- Polyhierarchical approaches can appeal for purposes, such as browsing, establishing connections within a body of knowledge and resolving internal disputes about where a concept 'belongs'.

- But they are unsuitable for financial reporting in which numbers must be allocated to a single 'bucket.' They can also become problematic to manage more generally.[4]

- We concluded that we should accept the constraint of monohierarchy, with expressivity being supported via our facets together with appropriate schemas, i.e. the data fields which are captured for matters, documents, people and so on.[5]

- Organisations with the necessary software and skills may, of course, adopt elements of noslegal or extend it as part of a more complicated polyhierarchical or ontological approach if this makes sense for them.

[4] Some of these issues are succinctly discussed here.
[5] For schemas, see Diagram 2 below.

4.3 Faceted

The noslegal taxonomy has a number of facets covering particular topics. This allows for a simpler list of concepts, with complexity being expressed by combining concepts rather than creating new sub-concepts.

Example. "Real estate litigation" might, in a problematic taxonomy design, be expressed in one context as a sub-concept of "real estate" and in another as a sub-concept of "litigation". Choosing one over the other reduces value and increases the potential for confusion and bad data. Choosing both is obviously a bad idea. A faceted taxonomy can solve this conundrum, without the complication of polyhierarchy, by putting real estate in a different facet from litigation with "real estate litigation" being expressed by combining concepts from different facets. This sort of complexity is addressed by combining facets rather than by sub-division. A sensible choice of facets (enough, but not too many) allows concepts to be limited in number while also offering great expressivity.

To illustrate the power of a faceted, hierarchical approach, a matter involving arbitration in Singapore concerning a trade union dispute in the mining sector in Indonesia can be expressed as shown in the way shown in Diagram 2 below. This allows for powerful:

- **Reporting and analysis** e.g. aggregating numbers by process, sector, legal area or geography to understand spend, profitability, staffing needs etc.

- **Search** e.g. if someone is looking for a credential or knowledge on "Asia employment dispute natural resources" then a matter classified in more detail can be returned if the search is driven by this taxonomy.

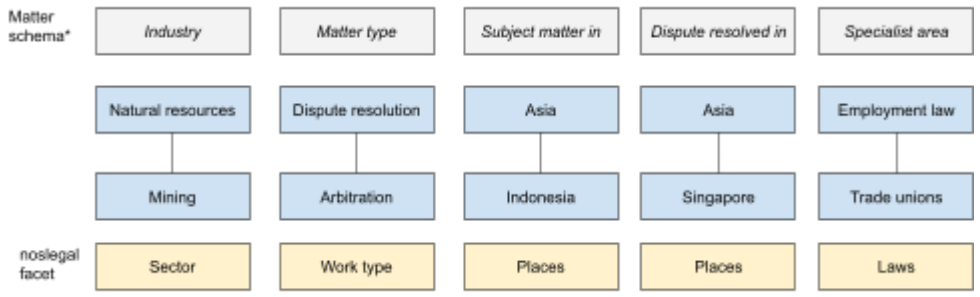

*'Schema' refers to the attributes which a matter (or client, supplier, individual, document etc) may have. The content may be drawn from a taxonomy or from a constrained list (often implemented as a dropdown) or be free text, numbers, currency and so on. Schemas and terminology vary between law firms and legal departments - the language used here is just an example.

Diagram 2

4.4 Limited size, modular and extensible

The noslegal taxonomy can be extended with further concepts for those who find it relevant. We believe this is the right strategy because:

- Limited concepts and levels are more practical to agree and maintain. Complexity, implementation obstacles, and the scope for differences of opinion increase as levels and the number of concepts increase.

- There is likely to be increasing AI-based application of taxonomies to data, and development of the finer details of taxonomies and ontologies. It is more useful to concentrate human effort on a fairly high level, high quality taxonomy.

- Simplicity also increases relevance across organisations, countries, types of legal work, and need. Each of these variables quickly leads to divergence the deeper into detail you go.

All that said, extensions can be very useful for niche needs:

- These can be produced as a joint effort within the noslegal community — we've already produced various extension packs, outlined in the release notes. Examining one of these will give an idea of what an extension looks like in practice. Please feel free to propose further packs[6] that you think would be useful to approach on a community basis.

- You're also free to produce extensions privately and either keep them private or publish them. Whatever you wish.

The modular approach (reflected in a combination of facets and core / extension packs within facets) also makes it easier to understand what's available and pick some elements while leaving others if that's most appropriate for your organisation. If doing that, do test whether our facets map appropriately to your approach, and consider with as open a mind as you can[7] whether you need to adjust one or the other to deal with any awkward partial fits.

4.5 Supported

A practical legal taxonomy benefits from explanatory materials and updates. We are committed to providing both. The release notes and other material (notably, a slide deck on practical implementation topics) which we have published since

[6] You can contact us via https://www.noslegal.org/

[7] In reality there are often constraints within a given organisation – concepts which 'cannot be touched' for whatever reason. Understanding iteratively what's possible within an organisation and what can be challenged can be a sensitive and trying process, but is worth pursuing where possible, building understanding between stakeholders in order to achieve something better together than any one group can achieve alone.

2022 can be found on GitHub. We are now publishing this guidance note, together with separate release notes describing the new taxonomy version.

Our ambition is to keep the existing facet cores as structurally stable as we can now, but there will be

- incremental improvements and updates,
- addition and refinement of more extension packs, and
- further facets when a real need for these is demonstrated.

5. Implementing the noslegal taxonomy

Once you've understood the merits of adopting a standardised taxonomy, you have the potentially tricky challenge of achieving buy-in within your organisation and implementing successfully in practice. This part addresses some of the main considerations.

5.1 Understanding the status quo

It's important to understand the full picture around data in your firm before taking any significant decisions. It is unlikely that all the data is "owned" by just one team. Mapping the flow of data of your organisation, within and between the four areas mentioned in section 2.2 of this guide, and their respective software applications and databases, is important for success.

Some of the questions you will need to ask are:

- Who are the stakeholders involved and who "owns" the various sets of data? See section 2.2 for examples of relevant teams. Coordinating between these is important, at a suitable level, so that higher-level concepts are shared or at least effectively mapped in ways that are usefully implemented in software.
- How mature and effective are existing taxonomies, processes, practices and software flows for classifying relevant data?
- Which software applications and databases process and store the relevant data? See section 2.2 for examples.
- To what extent, is there any integration between different data sources, and any existing attempts to combine it in separate tools (e.g. Tableau or MS Power BI)? What lessons can be learned from these?

- What demand, or at least openness, is there in each area to do better? Is there evidence of commitment to better data (e.g. existing manual efforts made to clean data)? Are the benefits understood already? Can they be assessed and to some extent quantified? Include the benefits of a coordinated approach between different areas in this exercise, as these can be huge in terms both of data usefulness and maintenance time/cost..

- What motivations are resonating with people? For example, the need to provide a taxonomical basis for application of AI to data has recently acquired increased prominence. This can be helpful, though be careful to ensure that the need to fix practical problems involving data is not diluted by a more general wish to 'use AI' (for example).

- Are those outside specialist functions, for example lawyers and senior managers, facing relevant problems which — whether they realise it or not — could be addressed by progress in this area?

In 2023 we published some slides which can assist with thinking about this.

5.2 Challenges

There tend to be three main barriers to implementation:

- **People and organisational realities.** If you're reading this guide you have probably already realised at least some of the merits of implementing a taxonomy. But convincing other people can be trickier.

    - Building a business case is important so that you can obtain top management's enthusiastic support. Successful change projects in this area are complex, and do need such support to make the necessary resources available, to communicate the importance and to address inappropriate refusals to engage or cooperate.

    - Helping individuals throughout the organisation understand how better data will benefit them can be challenging where data isn't at the forefront of their mind.

    - So think about how potential changes are communicated and how to win hearts and minds, by showing people how better data will help them and not be unduly burdensome on them to produce.

    - You may need to hold workshops, send out messaging and make sure senior leaders are on board by emphasising the clearest benefits and any particular problems which your organisation has encountered already (ideally expressed in financial terms and with

concrete examples, e.g. time spent, work lost, risks that have materialised).

- Have a short slide (a few bullet points or diagram) and a simple 'elevator pitch' which the implementation team members can repeat whenever there is a suitable opportunity. The risk of people not grasping the basic benefits is significant if they feel drowned in implementation detail.

- If you have particularly enthusiastic stakeholders, involve them in the discussion with senior management: a combined approach tends to be more persuasive than a particular specialised function seeking something on its own.

- It is important to get senior management commitment because, particularly in a legal organisation where people live and breathe abstract concepts, you will almost certainly run into objections from lawyers and others who are heavily committed to an existing way of conceptualising something and who may have unstated fears that a change in taxonomy will lead to some undesirable organisational change.

- Use concrete examples and data[8] to address demands which are likely to increase the risk of bad data (for example, by conceptual redundancy, unnecessary complication and burden on individuals).

- **Technology.** Most organisations are not starting from scratch when it comes to technology.

  - We have made a major effort to make noslegal implementable in a wide range of software through the design choices described earlier in this guide.

  - However, you may still run into limitations such as fields, characters, numbers of levels and more. In addressing these, it's important to put yourself in the shoes of the people who will be using relevant software, and making realistic decisions as to how much you can accurately capture using the software at your disposal, and the roles to be played by process, communication, training and broader

[8] For example, many implementation teams in law firms have found themselves forced to accept unsuitable concepts in the first implementation. A way of addressing this problem effectively is to review the classification data after a reasonable period (for example, a year) and seek permission to cull concepts which are hardly used, whether in financial terms (if used to classify matters) or in knowledge terms (number of documents or other items classified). In other words, demonstrate the problem objectively to gain permission to address it.

system considerations, not just the software. There is experience of this within the noslegal community so please feel free to engage.

- **Resource.** Really understanding how to get to grips with your data, and then actually doing it, takes some imagination, focus and a lot of time. Every organisation will have its challenges here.
    - Smaller organisations likely won't have a specialist taxonomist or data officer.
    - Larger organisations will have additional layers of organisational complexity and siloes to navigate.
    - The important point here is that you will need to put some significant investment into this, whether by diverting people from other work, or establishing a new role or using a consultant. There may also need to be some investment in new software or upgrades.

5.3 Pragmatism over procrastination

Despite the potential challenges, we advocate for doing something rather than nothing, and approaching the topic in stages. You will never improve your data outputs if you continue to be overwhelmed by the size of the task and simply ignore the issue.

- Projects can start small. Even if you are confined to a particular function within the organisation, think about where incremental changes could make the most impact in terms of having standardised terms, for example introducing naming conventions or additional metadata in documents. A high quality project in a particular team, with the benefits objectively assessed and communicated, can help to establish appetite elsewhere in the organisation by demonstrating feasibility and value.
- noslegal can be implemented as a standard with minimal change, but also feel free to amend it and to decide which modules to adopt or reject. Ensure though that you think through the need to make it all work together consistently, document it comprehensibly for people maintaining it in future, and have a strategy on how to update it effectively. Excessive customisation can be a trap otherwise.
- You don't have to tackle all your retrospective data at once. Making improvements means you will start building better data today, and after a few years much of the older data will often be less relevant.

- In terms of process and practices for applying a taxonomy to data such as matters, documents and credentials, “be the change you seek”. For example, if the process is not working well, figure out why and how to address this, talking to people and trying things out. This can be relatively low cost and is a great investment. You’ll need to do it anyway, sooner or later, or you’ll just be wasting the more substantial investments.

- Explore understanding of relevant topics across different functions informally over a period of time. It will probably take many months (possibly years!), but it builds a strong foundation. Problematic inconsistencies are often accidental, resulting from absences of communication.

- Test and iterate. When rolling out a new taxonomy, pay close attention to the numerous ways it can go wrong. Test the quality of the data, find out how people are feeling about it and identify the weaknesses - be they in taxonomy concepts, training, process,[9] software or anything else - so that they can be addressed. And in doing so, don’t underestimate the ability of better process to improve data quality even if the software you have isn’t ideal.

- Manage expectations as to what’s achievable in the short term, and try to deliver some fairly quick wins while making clear to people what won’t be possible immediately but how genuine needs can be addressed in future.

5.4 Further information

Many members of the noslegal community are open to sharing experience. There are also various consultants with experience in the field. Depending on demand and community members’ willingness to spend time on the work, we may issue further practical material on implementation in future. Please get in touch if you’re interested in this.

[9] Simply as an example of process issues, a major ‘bad data’ problem many law firms encounter arises from the fact that matters are classified upon opening, and in practice, the person filling in the matter-opening form may not know how to classify it. Also, the classification may change weeks or months down the line as the matter develops. Consider how, when and to whom you surface classification information to people in different contexts and how you can increase the accuracy. There are many possible tactics here and try to think imaginatively about them. Just as an example, expose only certain classification possibilities to people depending on the area of the firm they live in, though this will require careful work on what to make available to whom, and the user experience for invoking any exceptions. Another process example is, if you suspect that people are over-using a particular classification (that is, if it’s being treated as sort of ‘misc’ or ‘general’) then some organisations require an extra step for it to be used (for example, approval by a person who can be trusted to give the matter proper attention).

6. History and governance

noslegal was founded in 2020 as an informal group. We published a first version of our taxonomy in early 2022 and a second major version in 2023, with minor updates to the Places facet between these dates and in 2024. We established a company limited by guarantee early on to hold relevant IP and license it out on an open source basis. We have minimal operating expenses which are covered by modest sponsorship sums refreshed periodically (current names are shown on the website noslegal.org). In late 2024, with increasing adoption as illustrated in this September 2024 press release, we adopted the governance structure described in this document and summarised below:

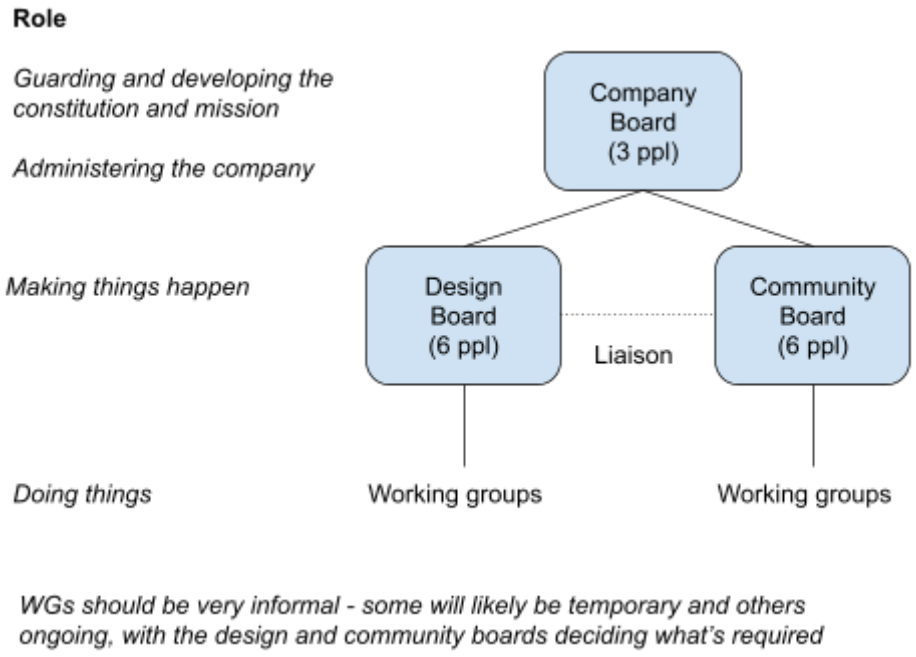

Diagram 3

Our Design Board meeting notes are published here and our Community Board meeting notes are here.