Charting the course with Kuberno:

Applied AI in Legal Entity Management

The Governance Guide

Kuberno
Navigating Governance

Table of contents

Section 1:	05	AI is forcing an operational rethink in governance	08
Section 2:	09	Agentic AI: opportunity or snake oil?	16
Section 3:	17	Finding the truth faster	20
Section 4:	21	Getting the work done	25
Section 5:	26	Staying ahead	29
Section 6:	30	What is coming and how to prepare	33
Section 7:	34	Ok, so what do you need to do first?	38

Charting the course with Kuberno: Applied AI in Legal Entity Management

Foreword

Over the past year, I have had more conversations about AI than almost anything else. Some have been exciting. Some have been practical. Quite a few have been slightly exhausting.

If you work in governance, company secretariat or entity management, you will probably recognise the feeling. Everyone wants to know your AI strategy. The board wants a point of view. The CFO wants productivity. The CIO wants assurance. Vendors want to show you agents. Meanwhile, your team is still dealing with filings, registers, approvals, board packs, statutory deadlines and the thousand small things that keep a business properly governed.

If this is your reality, then this guide might just be for you.

I don't want to argue about labels in this guide. Whether something is called AI, automation, workflow orchestration or an agent matters less than whether it makes governance better for your institution.

Does it reduce risk? Does it improve confidence? Does it give the business a clearer view of its entities? Does it free skilled people from repetitive work so they can spend more time on judgement?

That is the standard we should hold the technology to. However, caution should not become inertia.

AI is forcing us as a profession to zoom out and look again at how we operate. In my mind there is an underbelly to our profession that is shifting: from documents to data, and from static records to trusted operational workflows.

The organisations that understand this will build better governance functions. More connected. More auditable. More resilient. More able to give the business confidence.

That's why we have written this guide. A tool to decide what matters, what is credible, and what should come next in the AI journey we're all navigating.

For Kuberno, our view is simple. AI in entity management should be operated by technology, but controlled by humans. It should start with trusted data.

It should respect the risk profile of the function. And it should help governance teams do what they have always wanted to do: spend less time chasing records, and more time shaping the right strategic decisions.

We don't need to be first to everything. But we must strive to be right.



Zoe Bucknell
Kuberno Co-Founder and CEO

Executive summary

The future of AI in governance and entity management.

AI offers significant opportunities for governance, but in our view organisations should resist the temptation to chase the latest capabilities before the fundamentals are in place.

Being first to adopt the latest AI technology is unlikely to become a competitive advantage in itself. More often, success will come from being deliberate about where AI can create genuine value and taking the time to build the right foundations.

We think successful adoption tends to follow a fairly consistent pattern.

1

It starts with trusted data.

AI is only as reliable as the information it can access. If records are fragmented, inconsistent or difficult to trust, AI will simply expose those weaknesses at greater speed and scale. Getting "agent-ready" is likely to have far more to do with data quality than model selection.

2

Processes that are fit for an AI-enabled world.

Before asking what AI can automate, it's worth asking whether the underlying process is clear, consistent and well understood. In many organisations, AI is proving to be a catalyst for simplifying workflows, clarifying ownership and improving the way governance operates.

3

Choose the right tool for the job.

Not every problem needs the same technology. Some are best served by retrieval or a deterministic workflow. Others benefit from AI supporting particular steps, while some may justify more agentic execution.

The useful question is simply: what is the work, and what is the right tool for it?

4

Fundamentally, governance remains a human discipline.

However capable AI becomes, judgement, oversight, experience and empathy remain central to good governance. The opportunity is not to replace governance professionals, but to remove administrative burden so people can focus on the work where their expertise has the greatest impact.

Section 1:

AI is forcing an operational rethink in governance



Kuberno
Navigating Governance

Pressure to adopt, but less clarity on what to adopt.

Here is what your week probably looks like.

Last Tuesday, one vendor promised to "completely automate entity management". Another launched an AI certification programme. The CFO asked for an AI roadmap. A board member forwarded a LinkedIn post about autonomous agents. Meanwhile, the actual work of filings, resolutions, subsidiary tracking and compliance monitoring has not slowed down while your strategy catches up. If that sounds familiar, you are not alone.

Board-level disclosure of AI risk has tripled among Fortune 100 companies in a single year, while almost every CIO now reports increased pressure to demonstrate measurable returns from AI investment.

> **48% of Fortune 100 companies now cite AI in their board risk oversight disclosures, up from 16% a year earlier.**
>
> EY Center for Board Matters, 2025

Governance teams are under pressure to adopt AI without always having clarity on what will create the greatest value.

They are shown compelling visions of autonomous organisations and intelligent agents, yet the practical path from today's governance function to tomorrow's operating model often feels much less obvious.

That gap is becoming increasingly visible. Generative AI has found clear, valuable use cases in legal. More than half of in-house counsel now use it in some form, primarily for document review, legal research and summarisation.

> **52% of in-house counsel now use GenAI.**
>
> ACC/Everlaw GenAI report, 2025

> **Document review (77%), legal research (74%) and summarisation (74%) remain the leading AI use cases in legal. Only 11% of organisations currently have agentic AI systems in production.**
>
> Thomson Reuters, 2025 / Deloitte, Tech Trends 2026 – Agentic AI strategy

Governance presents a different challenge. Many of its highest-value activities extend beyond a single document or system.

Regulatory filings, cross-jurisdiction compliance, entity changes and governance workflows require AI to understand context, navigate multiple systems and maintain a complete audit trail. These are fundamentally different problems from summarising a document.

This is perhaps why so many governance teams feel caught between promise and practicality. Beyond the obvious generative AI use cases, the conversation quickly turns to agents. Yet for many organisations, the options either feel too immature for production, too risky for regulated environments or too disconnected from the way governance actually works. In our view, the conversation often starts with technology instead of the governance challenge.

A more useful starting point is to ask:

> **"Which AI capabilities are applicable to our governance challenges?"**

That requires an understanding of the work itself before selecting the technology.

A modus operandi shift. AI is forcing governance to rethink the work itself.

Generative AI has exposed the limits of digitising governance without fundamentally modernising it.

Over the last decade, paper became PDFs, wet signatures became e-signatures and boardrooms became Teams calls. But beneath those changes, the operating model remained largely the same.

Many governance processes are still organised around documents: forms, resolutions, filing packs and digital folders. AI is exposing a deeper reality.

Governance has always been driven by data, decisions and regulatory events. Documents are simply the record of those activities.

Take a director appointment. In many organisations it is still treated as a collection of documents to prepare, approve, sign and file. In reality, those documents are only the outputs of a broader governance process involving:

- Collecting and validating director information
- Identity verification
- Sanctions and disqualification checks
- Approvals and delegated authority
- Regulatory filings
- Updates to the entity records
- Changes to permissions, registers, board composition and group reporting

> **“Most governance work is not document work at all. A director appointment isn’t fundamentally a form. It’s an identity check, a regulatory submission and an authoritative data update.”**

Governance operations will require different AI capabilities. Some benefit from retrieval, others from workflow automation or decision support, while a smaller number may ultimately justify more autonomous execution.

A mindset shift.

We believe succeeding in the AI era requires governance to move from thinking in documents to thinking in systems. Governance is increasingly becoming a series of verified data events, permissioned actions and auditable state changes.

What does that look like?

A **verified data event** might be confirming a director's identity and accepting their details into the entity record.

A **permissioned action** might be automatically preparing statutory filings once the required approvals have been completed.

An **auditable state change** might be moving an entity from pending to filed, with every approval, submission and supporting record captured automatically.

> **Data quality is such a critical foundation for AI adoption. AI builds on the quality of existing data and processes. Where those foundations are weak, the value of AI will be limited.**

Our view is that organisations which continue to treat governance primarily as document administration will capture only a fraction of AI's potential. Those that rethink governance as structured data, connected workflows and auditable processes will be far better placed to apply the right AI capabilities, in the right places, at the right time.

Section 1 takeaways

- **Being first matters less than being ready.** AI delivers the greatest value when organisations have strong data, clear processes and a clear understanding of where different AI capabilities can help.
- **Governance needs an applied AI strategy.** Different governance challenges call for different capabilities, from retrieval and deterministic workflows to AI-assisted workflows and more agentic execution.
- **The next stage of digital transformation is operational.** Governance is evolving from document-centred administration towards structured data, connected workflows and auditable events.

Section 2:

Agentic AI: opportunity or snake oil?



Kuberno
Navigating Governance

“Agentic” has quickly become one of the most used terms in enterprise technology.

If you are working out an AI strategy for your governance team, there is a good chance you are already being shown agents by vendors. That creates a fairly practical problem: understanding what the technology can genuinely do, what level of autonomy is useful, and what you are actually buying.

We think agentic AI could be transformational in the right parts of governance. But it is more useful to think about ‘agentic’ as a spectrum of capability.

Some systems have very little discretion over how work gets done. Others can interpret what they are being asked to do, decide what happens next and adapt when something changes.

Understanding where a system sits on that spectrum makes it much easier to judge whether that extra autonomy is actually useful for the work you need it to do.

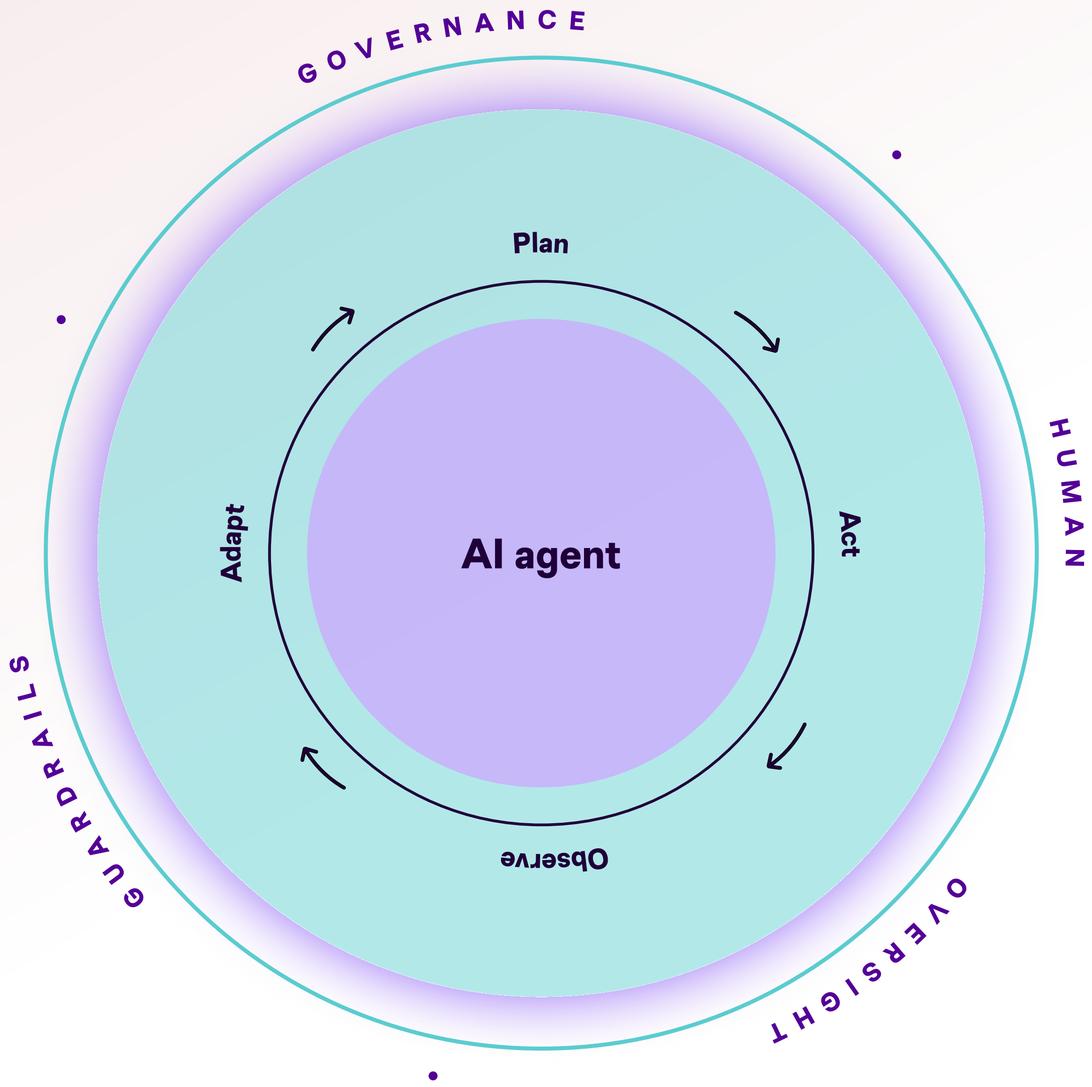

What does “agentic” actually mean?

The term is still evolving. Andrew Ng helped bring ”agentic” into the mainstream AI conversation in 2024, using it to describe systems with different degrees of agency. That is a useful way to think about it in governance too.

A workflow follows a route someone has designed in advance. As you move further along the spectrum, AI is given more freedom to interpret information, decide what to do next and respond when something changes.[1]

In practice, most agentic systems tend to do four things:

Work towards a goal

The outcome is clear, even if every step required to reach it is not.

Decide what happens next

They can select the steps or tools needed to move towards the goal.

Gather information as they go

They can draw on approved systems, records, documents and other sources as the work develops.

Adapt when the situation changes

They can ask for more information, try another route or involve a person when needed.

[1]Andrew Ng, ‘Welcoming Diverse Approaches Keeps Machine Learning Strong’, The Batch, DeepLearning.AI, 2024.

A spectrum of agentic autonomy in workflows

	DETERMINISTIC WORKFLOW	AI-ASSISTED WORKFLOW	AGENTIC EXECUTION
HOW IT WORKS	Route, steps and logic are defined in advance.	The overall process is defined, with AI supporting or completing particular steps.	A goal is defined, with the system given more discretion over how to achieve it.
WHO DECIDES THE ROUTE?	The process designer.	Mostly the process designer, with AI making limited decisions within defined steps.	The system can decide which steps, tools or route to use within its permissions.
WHAT HAPPENS WHEN SOMETHING CHANGES	Follows a predefined exception path or stops for human input.	AI may interpret the issue or request missing information while remaining within the workflow.	Can adapt, gather more information, choose another route or involve a person.
WHERE IT WORKS BEST	Repeatable, predictable processes with clear rules.	Processes that are broadly predictable but benefit from interpretation, extraction or generation.	Work where the route cannot be fully predicted in advance and adaptation creates value.

What does a real AI agent look like today?

To see where governance may be heading, let's look at legal, tax and finance to show what happens when AI starts taking more responsibility for how work gets done.



Legal: research and diligence

Legal is a useful comparison because the work is complex, document-heavy and often requires the system to plan its way through a task rather than follow one fixed route.

Harvey is a platform that can break a diligence exercise into tasks, gather information across documents, ask clarifying questions and return a structured review with supporting sources.[2]

Legora takes a similar approach to M&A diligence, reviewing thousands of data room documents for potential red flags, with lawyers remaining in control throughout.[3]

The difference is the shape of the work: diligence is often project-based, while entity management is recurring, operational and deadline-driven.

[2] Harvey, AI for Transactional Lawyers: Due Diligence & M&A, 2026.

[3] Legora, 'AI solutions for transactional lawyers', legora.com, accessed October 2026.



Tax: recurring filings

Tax is another close comparison: recurring, jurisdictional and deadline-driven, with approval before submission.

Avalara uses purpose-built agents across onboarding, data validation, filing, payments and notice resolution.[4] Its returns agent can ingest transaction data, apply the right forms and jurisdictional logic, and prepare work for approval.[5]

[4] Avalara, 'Agentic tax and compliance', avalara.com, accessed October 2026.

[5] Avalara, 'Avalara Launches Agentic Tax and Compliance™: AI Agents That Work for You', 30 September 2025.



Finance: period-end close

Finance is useful because period-end close brings together multiple systems, checks and specialist tasks around one defined outcome.

Sage Intacct's 2026 Release 1 added a new Finance Intelligence agent, working alongside four existing agents across areas including close, accounts payable, time and assurance.[6]

Together, they show how agentic systems can take on more of the work around a defined financial process without removing the need for human oversight.
That gets to the heart of the governance challenge: greater speed only creates value if people can still understand what happened, see the decisions made and step in where needed.[7]

[6] Sage, Sage Intacct 2026 R1 Release Notes, 13 February 2026.

[7] Sage, Sage announces Finance Intelligence Agent to power high-performance finance teams, 4 November 2025.

“One mistake I see repeatedly is governance teams waiting for the technology to become perfect before engaging with it. Every governance process today relies on judgement, interpretation and human decision making. I would not be asking whether an AI system or strategy is perfect before engaging. The real question is whether it can perform at the level of a competent professional while creating a stronger audit trail than the processes many organisations rely on today. The teams making the most progress are not necessarily buying the most AI. They are investing in the data foundations, operational discipline and workflow design that will allow them to adopt it safely when the technology is ready. Something I would stress again and again is that the barrier to agentic AI in governance isn’t the technology, it’s whether the underlying governance operation is ready for it.”



Stewart Foxall
Kuberno Chief Product Officer

When a vendor pitches an agentic capability, three simple questions to help you cut through the marketing:

1. How much of the route can the system decide for itself? A deterministic workflow follows a route designed in advance. Further along the agentic spectrum, the system has more freedom to decide which steps or tools it needs. Ask the vendor to show you what the system decides, rather than simply what AI is involved.

2. What happens when something unexpected occurs? Does the process stop, follow a predefined exception path or work out another way forward? Can it ask for missing information or bring a person in when needed? Ask to see an exception case, not just the happy path.

3. How is the AI itself governed? AI now has a formal international standard of its own. ISO/IEC 42001 covers how organisations manage AI risk, human oversight, transparency and ongoing monitoring, and can be independently audited and certified.

So another useful question to ask a vendor is **what standards do you use to govern your own AI, and can you evidence them?**

Practitioner POV: Zoe Bucknell, CEO, Kuberno.

> The question of auditing the work of an AI agent comes up often. My answer is that agents are capable of producing something humans find it hard to do: a complete, transparent record of every step taken and every decision made along the way. What information it used, what it decided, what it did, in what order. We will see this in production in entity management sooner than most people expect. So the question for governance teams is no longer whether AI can do auditable work. It's how you want to adopt it: what kind of audit trail your function needs, what your regulators and boards will want to see, and where you want humans positioned in the decision loops the agent runs through. Those are governance decisions, and they belong to us, not to the technology, or to the regulators.



Zoe Bucknell
Kuberno Co-Founder and CEO

Tips for procurement: The label matters less than the impact.

Procurement should start with the governance problem. A deterministic workflow may be enough. Another task may benefit from AI within the process. A smaller number will justify more agentic execution. The level of autonomy should follow the work. The next three sections look at where those opportunities are emerging across entity management: finding the truth, getting the work done and staying ahead.

Section 2 takeaways

- **Think of agentic capability as a spectrum.** Different systems give technology different levels of freedom to decide what happens next.
- **More autonomy does not automatically mean more value.** Use the level of agency the work actually needs.
- **Look at what the system does.** Can it interpret intent, choose a route, adapt when something changes and explain what happened?
- **More autonomy needs more governance.** Auditability, permissions and human oversight become more important as the system is given more freedom.

What AI is delivering for entity management.

Entity management involves three broad areas of activity:

Answering questions, executing work and anticipating risk. Each places different demands on AI, and each is progressing at a different pace.

Finding the truth
Reliable access to trusted information.

Getting the work done
Accurate execution across data, documents, approvals and systems.

Staying ahead
Continuous monitoring, interpretation and informed action.

Taken together, these areas help governance teams ask more useful questions about AI:

What kind of work is this?

Which AI capability is best suited to it?

What can be used safely today?

This provides a clearer basis for an applied AI strategy rooted in real governance needs. The next three sections look at each area in turn: where AI can already create value, where more advanced capabilities are emerging and what governance teams can do now to prepare.

Section 3:

Finding the truth faster



Kuberno
Navigating Governance

Finding the truth faster, the work.

Anyone running an entity management function will recognise how much time disappears into answering questions and repackaging information for different audiences.

- Who are the directors of our Luxembourg entities, and when were they appointed?
- What did the board approve in March 2021?
- Can you produce a structure chart for the refinancing?
- Where are the constitutional documents the bank has requested?

The information usually exists, but it is spread across registers, minute books, filing systems, document stores and jurisdictions, making even a straightforward request slow to answer and difficult for someone else to verify.



Does it need an agent?

Usually, no. Most of this work needs reliable natural-language search across the entity record, with clear citations showing where each answer came from.

As Section 2 established, this is fundamentally retrieval: a broader AI assistant, including an agentic one, may handle the task, but here it is simply being asked to find and present the right information.

Done well, the operational value is significant. Requests that once took hours or days can be answered in minutes. Teams spend less time searching and reconciling information, while reviewers receive both the answer and the evidence needed to trust it.

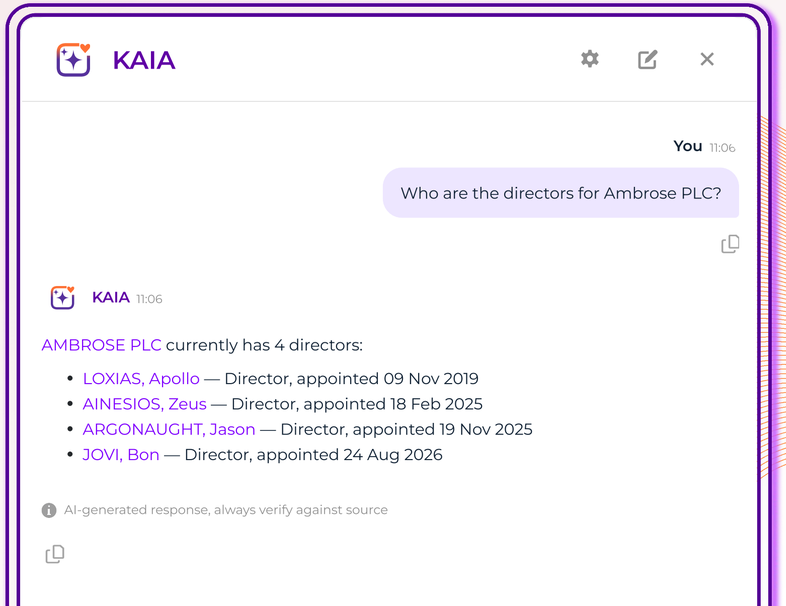

A practical use case

Consider a bank carrying out a periodic KYC refresh. It requests current directorships, shareholdings and constitutional documents for fourteen entities across six jurisdictions.

Today, that can easily mean several days of checking registers, searching through document stores and following email chains to work out whether the information is current. With natural-language search across a clean entity record, the same request becomes a query. The system can assemble the information, produce an exportable pack and provide a citation trail that a reviewer can verify in minutes. The same capability can support structure charts, board history queries, transaction packs, audit requests and regulatory information exercises.

What exists and what is coming

Search and retrieval are currently the most mature AI capabilities in entity management. Tools can already search across entity records and supporting documents, summarise information and extract relevant details. Increasingly, they can present that information in a format suited to a particular audience, including transaction-specific dashboards, KYC packs and structure charts.

Kuberno's own approach reflects this direction, with our AI Assistant - KAIA. You can ask natural-language questions about your entities and officers and receive practical outputs, such as structure charts tailored to a specific request or transaction.

The threshold should be one that is at least as accurate as one produced by a competent company secretary under time pressure, with every material point linked back to its source. Without that evidence, the team still has to repeat much of the original work to verify the result. Worse, an answer that sounds plausible may be relied upon without anyone spotting that it came from an outdated document, an incomplete record or the wrong entity. Speed only creates value when the answer can be checked.

What teams should do now

Build an entity-centred source of truth. Search will only ever be as reliable as the records beneath it. Fragmented data, duplicated documents and unclear ownership will produce incomplete or misleading answers, however impressive the technology appears in a demonstration. In practice, that means:

Identifying the authoritative source for each entity data point.
Teams should know whether a director's address, appointment date or ownership record is controlled by the entity management system, an external register or another approved source.

Resolving duplication and inconsistency.
Different spellings of a director's name, conflicting appointment dates or several versions of the same constitutional document need to be reconciled before AI is asked to use them.

Connecting documents to the events they evidence.
A board resolution should be linked to the appointment it authorised, a filing receipt to the record it changed and a constitutional document to the entity and period for which it applied.

Defining permissions around sensitive information.
The access controls already governing personal, legal and commercially sensitive data need to carry through into any AI search or assistant.

Establishing clear ownership for data quality.
A named person or team should be responsible for resolving errors, approving changes and maintaining the integrity of the record over time.

As we've reiterated, quality data is the first foundation for useful AI. It also creates value before any new capability is introduced. Records become easier to trust, ownership becomes clearer and teams spend less time reconciling competing versions of the truth.

Section 3 takeaways

- **Finding the truth is the most mature operational AI opportunity in entity management.** Natural-language search can already reduce the time spent locating, checking, assembling and reformatting information.
- **Reliable answers require a reliable entity record.** Search quality depends on accurate data, consistent naming, current documents, clear ownership and an agreed source of truth for each material data point.
- **Every material answer should cite its source.** A reviewer should be able to trace a statement back to the exact record or document, see when it was updated and confirm that it applies to the right entity.
- **Data preparation creates value now.** Cleaner records and clearer ownership improve everyday governance while preparing the function for more advanced AI.

Section 4:

Getting the work done



Kuberno
Navigating Governance

Getting the work done.

A significant part of entity management begins after a decision has already been made.

The board may have approved a director appointment, an ownership transfer may have been agreed or a statutory deadline may be approaching. The judgement and authority behind the change already exist, but the work required to implement it is far from complete.

Information still needs to be checked, approvals recorded, documents prepared, registers updated and submissions made across the relevant systems and jurisdictions. Annual filings, director changes, ownership updates and other recurring statutory activities therefore involve far more than the final form. They require accurate, auditable execution across several stages, often involving multiple teams and platforms.



Does it need an agent?

This work will probably span the full spectrum: deterministic workflows, AI-assisted workflows and, in some cases, more agentic execution.

Many processes are still best suited to a deterministic workflow, because the route is already known.

A routine annual filing, for example, could begin with the system identifying the approaching deadline, retrieving the latest entity data, comparing it with the previous submission and flagging any changes. It could then prepare the filing for review, record the approval and update the entity record once the submission has been accepted.

In that situation, the steps and approval points are defined in advance. People remain responsible for decisions carrying legal or statutory significance, while technology extracts information, validates records, identifies inconsistencies and prepares the required documents.

Greater agency becomes useful where the process needs some interpretation or flexibility.

At the start, that might mean understanding an instruction such as ”appoint Jane Smith as a director” and working out which workflow, information and approvals are needed.

It can also help when something goes wrong. A registrar may reject a filing, a mismatch may appear in the entity data or a required document may be missing. A more agentic system can interpret what has happened and decide whether to continue, try another route or involve a person.

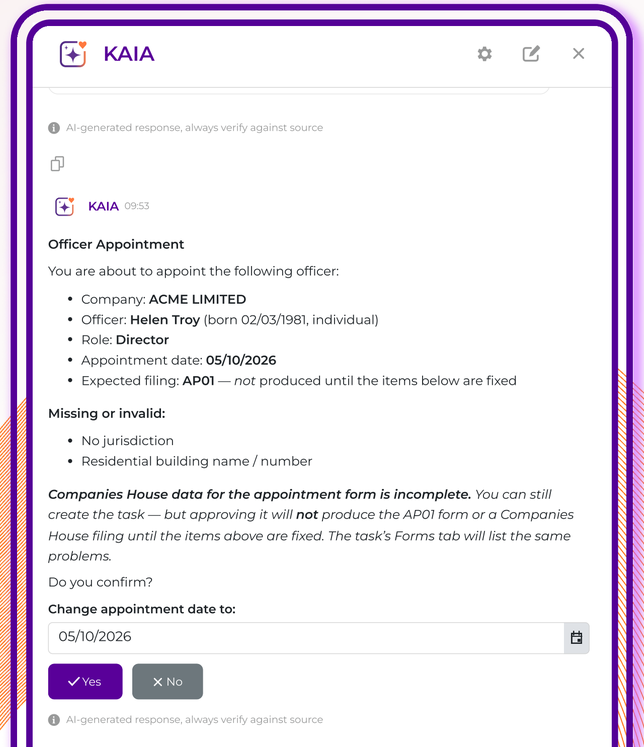

A practical use case

Take a director appointment.

The system collects and validates the director’s details, triggers identity verification and sanctions checks and records the relevant approval. It then prepares the statutory forms and internal documents, queues the filing for review and updates the registers once confirmation is received.

Human approval remains at the points carrying judgement, authority or statutory risk. That may include approving the appointment, resolving a discrepancy in the director’s information or authorising the final filing. If the registrar rejects the submission, an agent can assess the response. It may correct a formatting issue and prepare a revised filing, or escalate the matter to the governance team with the rejection reason, supporting information and suggested next step.

The process remains controlled, while the need for a person to move and re-enter the same information manually at every stage is reduced.

What exists and what is coming

A growing amount of execution work can already be supported by technology.

Entity data can be drawn from the entity record, checked and used to populate forms and documents.

Kuberno's workflow model already sits across parts of this spectrum. Known governance processes can move through predefined documents, approvals and downstream updates, while AI can support particular steps.

Kube can also understand the intent behind a natural-language instruction and identify the process needed to act on it, giving some workflows more flexibility than a fixed rules-only route.

A director change, for example, can trigger the relevant resolutions, statutory forms and approvals before updating the entity record once the change is confirmed.

If a filing is rejected, it might interpret the reason, determine whether the issue can be corrected automatically and decide whether human review is needed. This points towards a more connected form of execution, where an authorised change can move through the entity management platform, approval system, internal registers and external filing environment without the same information being repeatedly entered and reconciled.

The pace of execution will also vary considerably by jurisdiction. Registrars themselves sit at very different levels of digital maturity: from paper and wet-ink requirements, through email and portals requiring manual keying, to structured submissions and, at the most advanced end, open APIs.

The same governance process may therefore move almost automatically in one jurisdiction and still require significant human handling in another. For multinational teams, another challenge is simply having visibility of those differences, particularly where local rules sit with individual providers or are buried in historic process documents.

Optimise before you automate

Preparing for automation involves more than documenting the process currently in place.

From the governance teams and businesses we have worked with, the official process is rarely the whole picture. Years of local workarounds, duplicated controls, spreadsheets and manual handovers become embedded in how the work gets done. Some remain necessary. Others survive because no one has had the time or authority to challenge them.

Automating that process without first improving it simply carries the same complexity into a faster system.

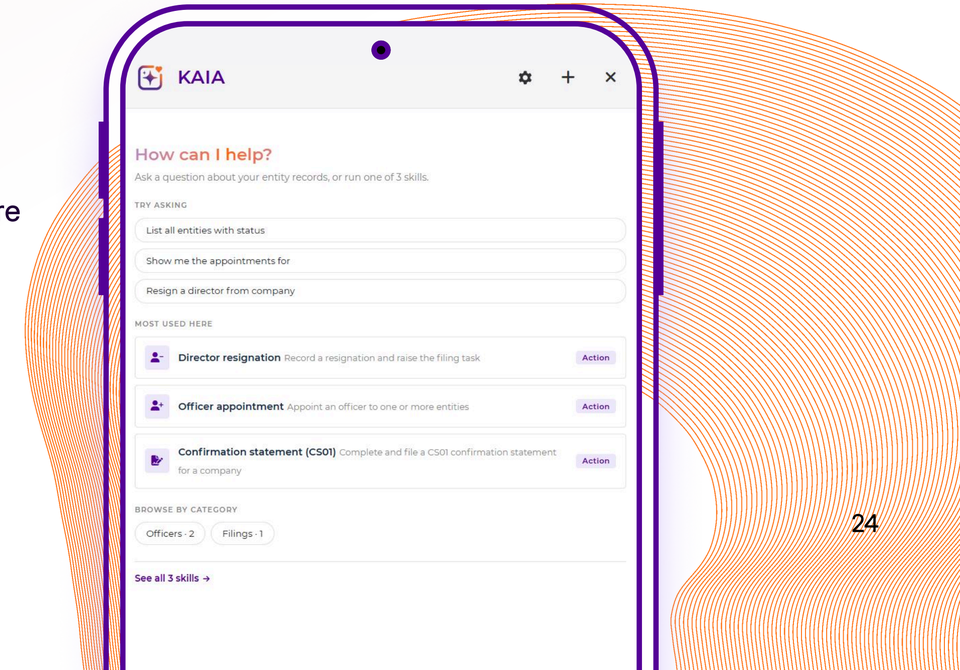

A useful discipline is to work through the process in stages.

Map it	Document what actually happens, including informal workarounds, duplicated checks, spreadsheets and manual handovers.
Challenge every step	Ask why each activity, approval and data entry exists, who needs it and whether it still serves a legal, regulatory or operational purpose.
Remove what is unnecessary	Eliminate duplication, legacy controls and steps that add little value. If testing does not require some elements to be added back, the process may not have been challenged hard enough.
Redesign what remains	Apply workflows and AI once the redesigned process is stable, understood and capable of being monitored
Automate last	Create the simplest process that still delivers the required governance outcome, with clear ownership, authoritative data and approval gates proportionate to risk.

Section 4 takeaways

- **Execution sits across a spectrum.** Deterministic workflows, AI-assisted workflows and agentic execution each have a role depending on how predictable the work is and how much judgement it requires.
- **Greater agency is useful where interpretation and adaptation matter.** Understanding intent, responding to exceptions and deciding what happens next are the behaviours to look for.
- **Process optimisation comes before automation.** Map the work, challenge every step, remove unnecessary activity and redesign the process before applying technology.
- **AI multiplies the operation beneath it.** Strong processes become faster and more scalable. Weak processes produce more errors and confusion.
- **Statutory execution depends on accurate data and clear control.** Well-structured information, proportionate approval gates and a complete audit trail are all essential.

Section 5:

Staying ahead

Kuberno
Navigating Governance

Staying ahead, the work.

Some of the most strategically valuable governance work begins before a request reaches the team.

It involves noticing what has changed, understanding what it means for the organisation and deciding what deserves attention.

Looking outward, that might involve identifying a regulatory change in one jurisdiction and mapping it to the entities affected. Looking inward, it could mean recognising that a new subsidiary has not entered the governance process, a shareholding threshold has been breached or a director resignation has created filing and board composition implications across several jurisdictions.

It may also involve identifying opportunities. An organisation could have dormant or duplicative entities suitable for rationalisation, but no one has had the time or group-wide visibility to spot them.

This kind of continuous monitoring can strengthen the strategic role of the company secretary.

It provides earlier insight into emerging risks, highlights opportunities for operational improvement and gives the board greater visibility of how changes across the group may affect its obligations and decisions.

How agentic does it need to be?

Staying ahead sits further along the agentic spectrum because the route is much harder to define in advance.

A goal such as "**make sure no entity in the group misses a regulatory obligation**" requires the system to monitor different sources, work out what is relevant, connect that change to the right entity data, gather anything that is missing and decide what should happen next.

That is where greater agency starts to become genuinely useful.

Used well, this capability could give company secretaries earlier and more complete visibility of issues developing across the group. That creates the opportunity to bring insight to the board before a risk becomes urgent, strengthening the vertical visibility and strategic contribution of the governance function.

The same capability also makes human judgement more important. Regulatory changes are rarely applied through data alone. They require interpretation, awareness of local context and an understanding of how legal requirements interact with the structure, priorities and risk appetite of the organisation.

As the system moves from identifying information towards interpreting its significance and recommending action, governance professionals remain responsible for deciding what is material, what response is appropriate and when the board needs to be involved.

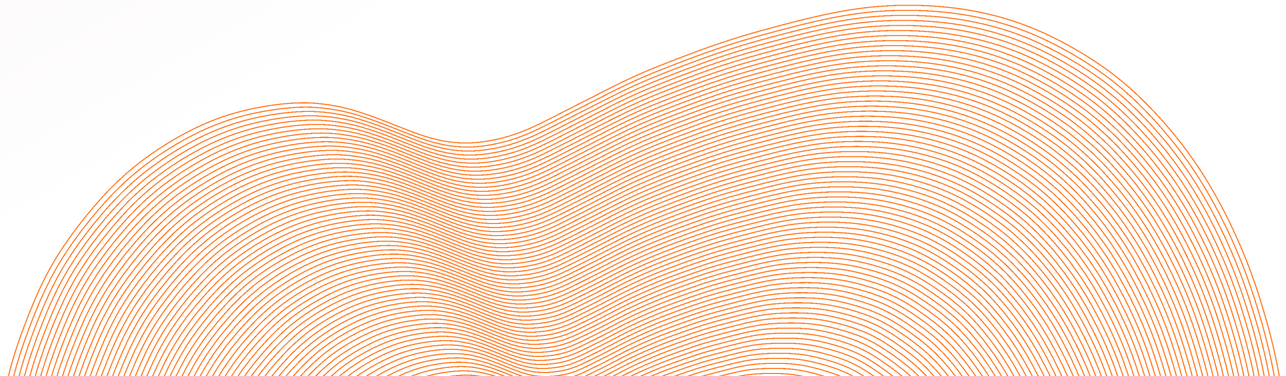

A near-future use case

Imagine a beneficial-ownership rule changing in one jurisdiction.

An agent monitoring trusted regulatory and legal sources identifies the change and assesses its relevance to the group. It determines that three entities may be affected, produces an impact summary and prepares remediation workflows.

The governance team receives:

- the source of the regulatory change
- the agent's interpretation
- the entities it believes are affected
- the data used to reach that conclusion
- the proposed actions
- the points requiring human judgement

The team reviews and approves the response before anything carrying statutory consequence is executed.

No one had to ask the initial question. The system noticed the change, connected it to the organisation and prepared a controlled response. The governance team could then concentrate on interpretation, materiality and the decision to act.

What exists and what is coming

Kuberno's **Registrar Intelligence** brings registrar and jurisdiction information into one place, giving teams a faster way to understand the rules before they file. It covers registrar requirements, governing company law, entity and filing types, register statuses and appointment rules across the jurisdictions an organisation operates in.

Teams can explore that information through an interactive map or structured directory, and narrow the view to the jurisdictions relevant to their group. In practice, that means less time searching old process documents, checking multiple registrar websites or going to local counsel for relatively basic jurisdictional questions.

It also points towards a more proactive model of governance, where external regulatory information can increasingly be connected with an organisation's own entity data to help teams understand which entities, structures or obligations may be affected by a change. Human judgement remains essential. The information is compiled from public registrar, government and legislative sources and should be verified against the official source or qualified counsel before it is relied upon for a filing or material decision.

What teams should do now

Governance teams can begin by defining the operating boundaries.

A useful framework is to separate what a system may:

1. **Flag.** Identify a potential issue, change or opportunity and bring it to a person's attention.
2. **Prepare.** Gather evidence, assess impact, draft recommendations and open the relevant workflow.
3. **Execute.** Take an approved action across internal or external systems.

The appropriate boundary will depend on the risk and reversibility of the action. Flagging a potential filing obligation carries a different level of risk from submitting the filing.

Teams should also decide:



Which sources the system may monitor. This could include official gazettes, regulators, company registries and approved legal information providers.



Which actions require explicit approval. External filings, changes to authoritative records and communications to the board should normally require a named person to approve them.



How authoritative those sources are. Primary legislation or regulator guidance may carry more weight than commentary, alerts or secondary summaries.



How interpretations and decisions will be recorded. The record should preserve the original source, the system's assessment, the reviewer's judgement and the final decision.



What confidence threshold is required. A low-confidence match might be flagged for review, while only a high-confidence result should be allowed to trigger a workflow.



How the system should respond when evidence conflicts. It may need to pause, present the competing sources and escalate the matter rather than selecting one autonomously.

These governance decisions can be made before the technology becomes widely available.

Section 5 takeaways

- **Proactive governance offers some of the greatest strategic value from agentic AI.** Continuous monitoring and interpretation can help governance teams identify risk earlier, find opportunities and provide stronger insight to the board.
- **Early capabilities are already emerging.** Regulatory intelligence, legal analysis and entity data are beginning to come together, although fully closed-loop execution remains less mature.
- **Human judgement becomes more important as the work becomes more proactive.** Governance professionals must interpret materiality, local context and the appropriate response.
- **Preparation begins with decision rights.** Teams should define what a system may flag, prepare and execute.
- **Auditability must grow with autonomy.** A proactive system should show what it observed, how it interpreted the information and why it recommended or took an action.

Section 6:

What is coming and how to prepare



Kuberno
Navigating Governance

1. Cross-system execution: the next real test

One of the most important developments over the next twelve to eighteen months will be AI working across the systems governance teams already use.

Entity management rarely happens in one place. A task may involve the entity record, board portal, approval system, document store, regulatory register and external filing environment. Connectors and standards such as Model Context Protocol are beginning to make those systems easier for AI to work across.

Kuberno's **MCP Connector** is already an example of that shift. Kube data can be accessed directly through Claude, Microsoft Copilot, ChatGPT and custom enterprise AI environments. Users can find entities, officers or specific data points in plain English, without navigating through Kube itself.

The important caveat is that this is not yet the same as fully autonomous cross-system execution. Today, the connector makes trusted Kube data available inside other AI environments; the next step is enabling AI to act across multiple systems as part of a single governed process.

Crucially, those interactions still respect Kube's existing record- and field-level permissions. For governance teams, that is the standard cross-system AI needs to meet: trusted data, existing access controls and a clear line back to the entity record.

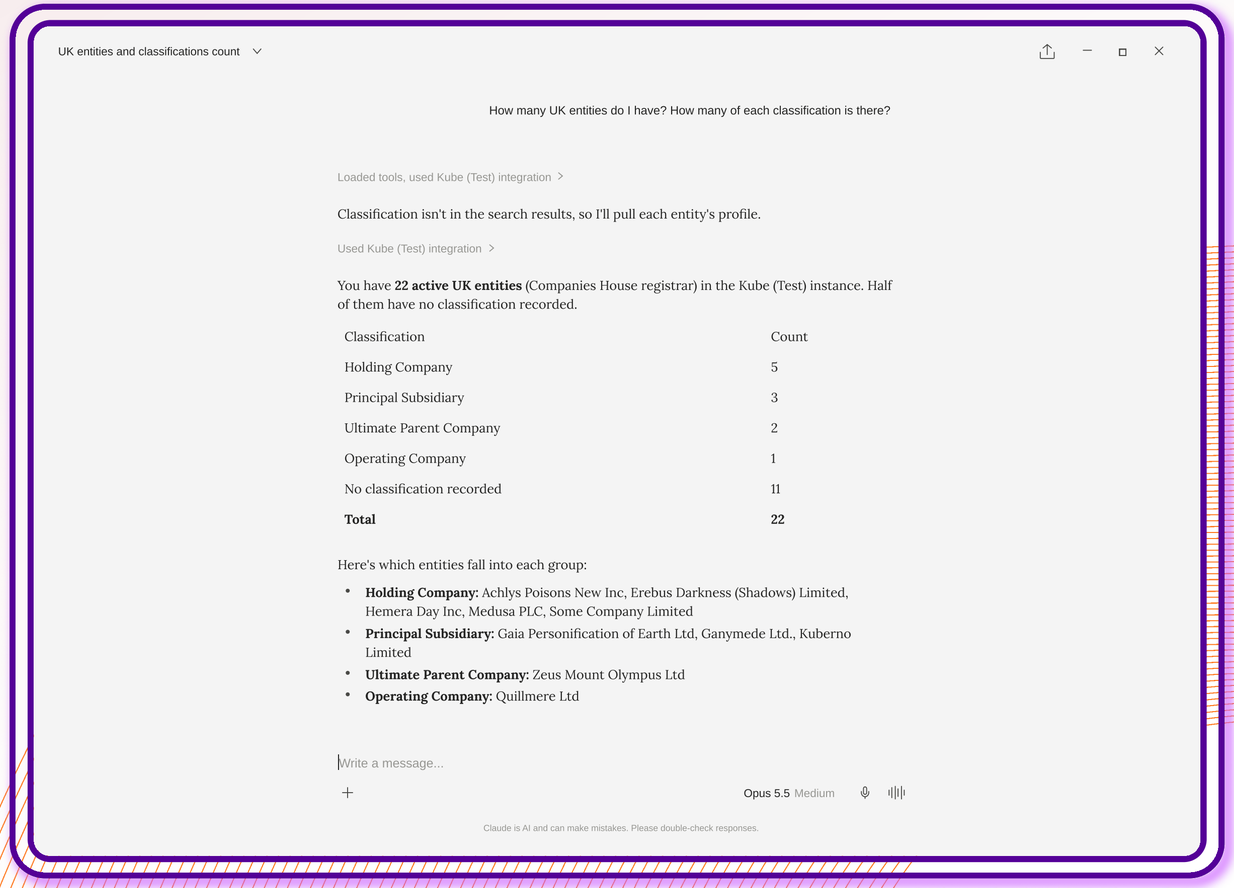

2. Create a safe environment for learning

Governance teams do not need reminding why caution matters. Entity records contain sensitive corporate, personal and regulatory information, and the consequences of getting things wrong are real. But secure enterprise environments now make it possible to experiment more safely.

Start small. Use dummy or non-sensitive data, choose lower-risk use cases and be clear about what you are testing. Natural-language search, duplicate detection or draft structure charts can all help teams understand how the technology behaves and where the real limitations sit — whether in the AI, the process or, often, the underlying data.

A simple principle we'd recommend is to experiment where failure can be contained, and apply a much higher bar where statutory, regulatory or data-security consequences are involved.

Critical use cases require a higher standard and a more controlled route to adoption.

Agents supporting statutory compliance, regulatory interpretation or external filings need robust architecture, specialist knowledge, secure integrations, tested controls and clear accountability.

They also require ongoing maintenance as regulations, registries and organisational structures change. Trusted specialist providers have an important role in developing and maintaining the infrastructure needed for this high-value, high-risk work.

Governance teams should assess those providers against real scenarios, including exceptions and failures, and require evidence of how permissions, approvals, data quality and auditability are managed.

This approach allows teams to build direct experience where the consequences are contained, while applying a much higher bar to systems carrying statutory responsibility.

Leadership's role

Leadership should make experimentation safe, purposeful and useful. Each pilot needs a clear problem to solve, defined data and permissions, agreed points for human review, and someone accountable for the outcome.

The important part is what happens afterwards. Teams should be open about weak results, unexpected behaviour and failed assumptions, then decide whether to stop, refine or scale the experiment.

Careful adoption still involves failure. The goal is to make sure it happens in the right places, with limited consequences and clear learning.

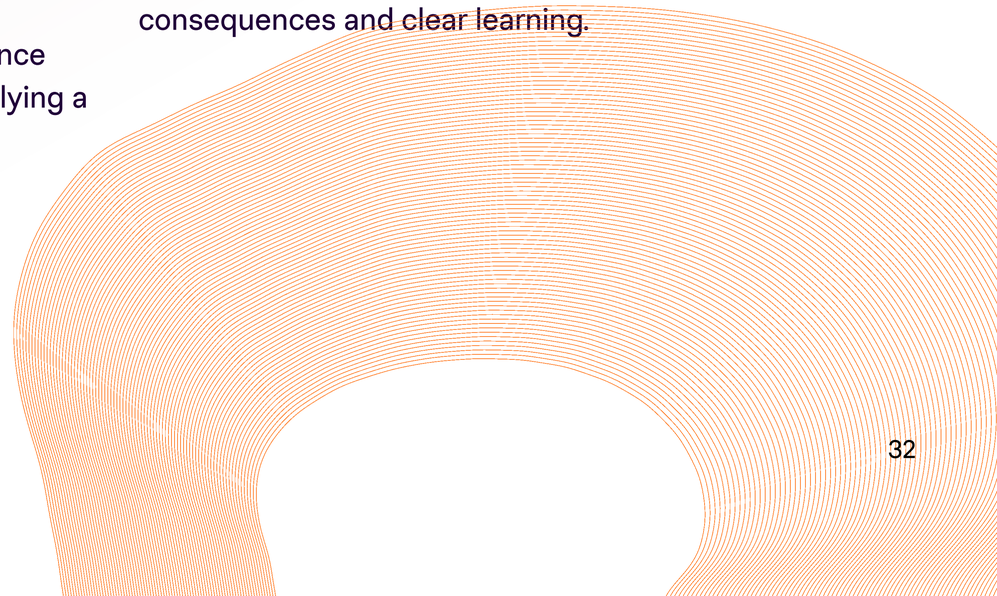

3. Prepare for an uneven world

AI adoption will develop unevenly across jurisdictions.

Some registries and regulators will move quickly towards structured data, automated submissions and AI-enabled interaction. Others will remain reliant on PDFs, portals and manual processes for much longer. A multinational organisation may soon operate across both worlds at once.

One jurisdiction may accept structured data through an automated interface. Another may still require a locally formatted form and wet-ink supporting document.

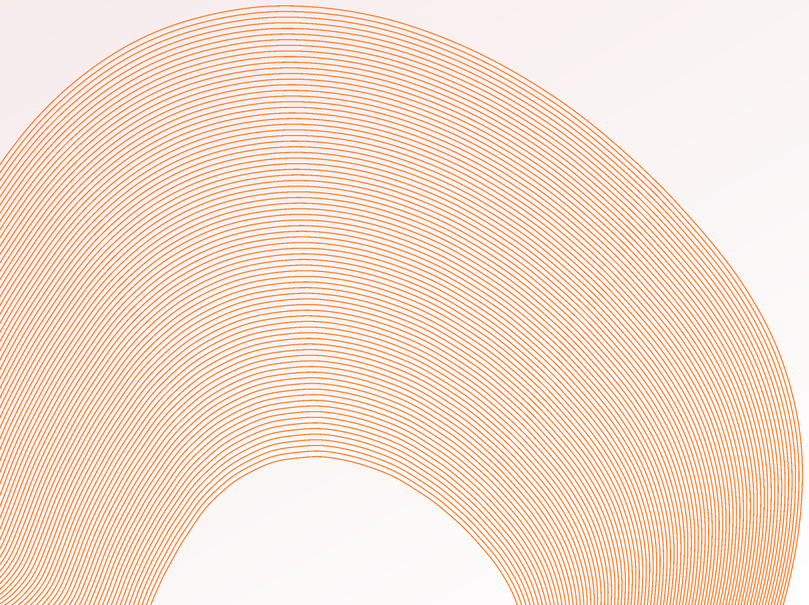

Build capabilities that travel

The most durable preparation is the development of skills that remain valuable regardless of which platform or AI model succeeds.

Those skills include:

- Process mapping and optimisation
- Data stewardship
- Technology fluency
- AI and vendor evaluation
- Risk assessment
- Deciding where human judgement and approval sit

Governance professionals who can explain how work gets done, identify the entity data it depends on and decide what a system may flag, prepare or execute will be much better placed to evaluate whatever technology arrives next.

Technical fluency does not mean every company secretary needs to become an AI engineer. It means understanding enough to ask good questions, challenge a vendor and make informed decisions about where the technology belongs.

Section 6 takeaways

- **Cross-system execution is the development to watch.** The practical value of AI will increasingly depend on its ability to operate safely across records, approvals, portals and filing environments.
- **Experiment where failure can be contained.** Approved enterprise environments, dummy data and lower-risk use cases help teams build experience without placing statutory obligations at risk.
- **Treat experimentation as a governance process.** Every pilot needs a clear purpose, defined data, limits on action, human checkpoints, measurable outcomes and an accountable owner.
- **Apply a higher bar to critical execution. High-risk** and high-value use cases require trusted specialists, robust architecture, strong controls and clear human accountability.
- **Adoption will remain uneven.** Governance operations need enough flexibility to work across jurisdictions and systems at different levels of digital maturity.
- **The strongest preparation is transferable capability.** Process optimisation, data stewardship, technology fluency and knowing where human judgement belongs will matter whatever tools emerge next.

Section 7:

OK, so what do you need to do first?



Kuberno
Navigating Governance

Can you trust your data?

Is your entity data accurate, current and structured? Do you know which system holds the entity record? Are duplicate or conflicting records being resolved?

Yes? Good. Move on. If not, start here. AI will only be as useful as the information beneath it.

Are your processes actually working?

Can you clearly explain how a director appointment, filing or entity change moves from start to finish? Are there duplicated checks, workarounds or manual steps that could be removed?

Yes? Move on. If not, optimise these processes before you automate them. AI will multiply a good process, but it will also multiply a bad one.

Do you know which work you want AI to help with?

Be specific about the problem. Are you trying to find information faster? Automate a repeatable workflow? Handle exceptions? Monitor regulatory change? Spot risk before someone asks?

Different problems need different capabilities. Start with the use case and work backwards to the technology.

Have you decided where people stay involved?

Be clear about what AI can **flag**, what it can prepare and what it can **execute.** Higher-risk actions should have stronger human oversight. Judgement, materiality, context and accountability remain a crucial part of the job.

Can you experiment safely?

Start somewhere useful but contained. Use approved environments, lower-risk use cases and dummy or non-sensitive data where appropriate. Give yourself permission to learn. The important thing is knowing where failure is acceptable, and where it absolutely is not.

OK, now you're ready to get started.

When you start talking to vendors, the product demo is only part of the story.

Can I trust the answer?

Can every important output be traced back to an authoritative source?

What happens when something goes wrong?

Does the system stop, adapt or involve a person?

Can I see what it did?

You should have a clear audit trail of the information used, decisions made and actions taken.

How are permissions handled?

Existing access controls should still apply when your data is being accessed through AI.

How agentic is it, and do you actually need that?

What can the system decide for itself? Can it interpret intent, choose a route, use different tools and adapt when something changes? Then ask whether those capabilities genuinely improve the use case.

Can it work across systems?

Look at how it connects to the tools your team already uses, including emerging standards such as MCP.

Does it understand the jurisdictions you operate in?

Entity management is local by nature. Regulatory and registrar intelligence needs to reflect that.

Where does human judgement sit?

The answer should be clear before anything carrying statutory or regulatory consequence is automated.

How do they govern their own AI?

Ask what sits behind the product as well as what you can see in the demo. ISO/IEC 42001 is a useful benchmark because it provides an independently auditable standard for how an organisation manages AI risk, human oversight, transparency and ongoing performance.

Where Kuberno can help

Kuberno is building around the same principles set out in this guide: trusted entity data, well-designed processes and AI applied where it can create practical value.

KAIA gives teams a simpler way to access the information already held in Kube, using natural-language questions to find entity and officer data and produce useful outputs such as structure charts.

Kube workflows help turn that trusted information into action, moving entity changes through the right documents, approvals and downstream updates in a controlled and auditable way.

Registrar Intelligence adds the external context, bringing registrar requirements, company law, filing types, register statuses and appointment rules together across the jurisdictions an organisation operates in.

And the **Kuberno MCP Connector** connects that trusted entity data to the wider AI tools teams are already using. Kube data can be accessed through Claude, Microsoft Copilot, ChatGPT and custom enterprise AI environments, while retaining existing record- and field-level permissions.

We hold ourselves to the standard. **Kuberno achieved ISO/IEC 42001 certification in March 2025** and has since completed its first annual surveillance audit with no issues or concerns raised.

If we are asking governance teams to look closely at how their AI vendors operate, we should be able to show the same level of scrutiny ourselves.

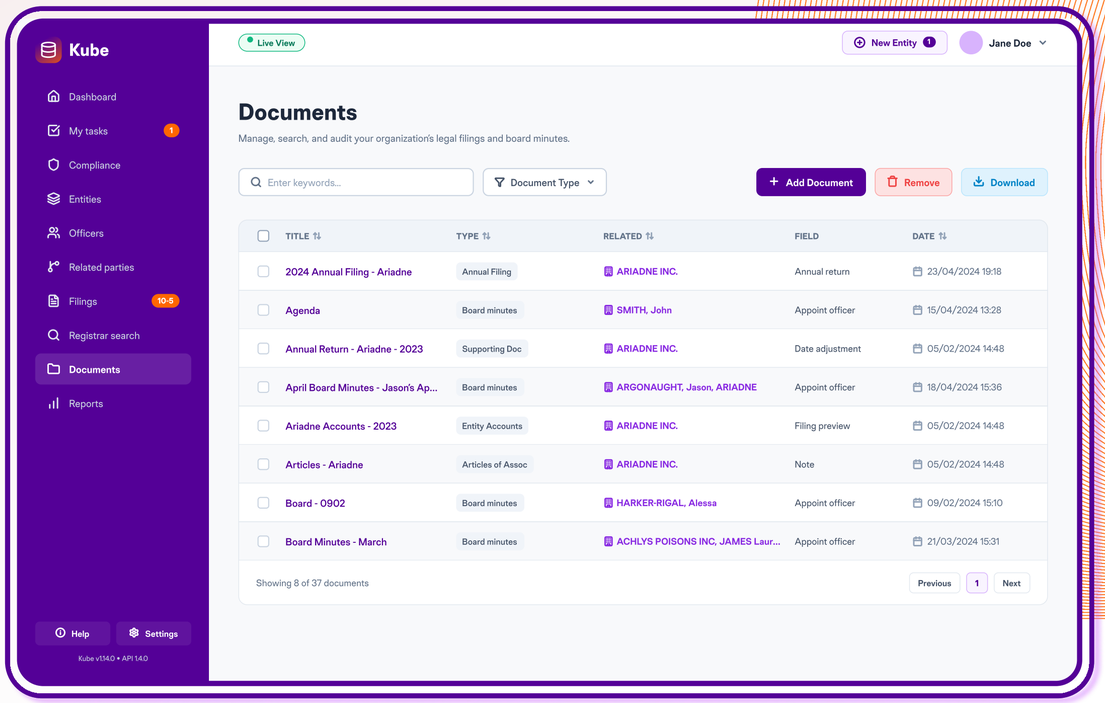

Taken together, we make it easier to find the truth, get the work done and stay ahead, without losing the controls and human judgement good governance depends on.

Want to work out where to start? Chat to a Kuberno consultant about your data, processes and AI priorities, and where the biggest practical opportunity sits for your team.

Charting the course with Kuberno:

Applied AI in Legal Entity Management

Thanks for joining the journey. If you want to explore how Kuberno can help your governance team move faster, reduce risk, and make better decisions, we'd love to talk.

Find out more →

Kuberno
Navigating Governance