

LEGAL AI DESIGN

**The Methodology for
Productizing Legal Expertise**

Dr. Gert Loterman



Copyright and Use

© 2026 Gert Loterman

All rights reserved. No part of this book may be reproduced, stored in a retrieval system, or transmitted in any form or by any means—electronic, mechanical, photocopying, recording, or otherwise—without the prior written permission of the copyright holder, except for brief quotations in reviews or scholarly works.

Executive Summary

Legal expertise traditionally resides primarily in individual lawyers, who repeatedly apply their knowledge, judgment, and methods to individual client matters. Even when similar matters recur, applying that expertise requires lawyer time, making that expertise difficult to reuse and scale independently of the practitioners who possess it. Traditional legal services therefore largely commercialize expertise through lawyer time, typically billed by the hour. The firm's capacity and revenue remain closely tied to the professional hours available, making this business model difficult to scale.

Productization addresses this problem by transforming legal expertise into reusable Legal Assets: Document Templates and AI Workflows. Document Templates capture drafting knowledge in predefined content and rules, while AI Workflows capture analytical methods in structured steps that respond to each matter's information. In productized legal services, these assets allow expertise to be applied repeatedly across matters with less recurring lawyer effort, supporting more standardized and scalable legal work, as well as more predictable pricing where the resulting services are offered to clients. Legal AI Design provides the methodology for making this transformation.

Foundations

A Legal AI System provides the environment for this transformation by connecting Case Management, Document Automation, and Legal Intelligence. Case Management organizes legal work in Client Cases, which bring together the information and work associated with individual client matters, and Knowledge Cases, which contain reusable legal knowledge for use across matters. Document Automation draws on Client Case information to generate matter-specific documents. Legal Intelligence supports contextual analysis and language-based legal work using information from both kinds of Case, and can also help develop reusable assets. Retaining the resulting work and assets within this environment allows expertise developed in individual matters to be refined and reused. Document Templates and AI Workflows constitute the two kinds of Legal Asset addressed by Legal AI Design.

Document Templates allow established drafting decisions to be applied consistently across matters. A precedent preserves earlier wording; a Document Template also captures the decisions governing how that wording should change to suit the matter. Where wording and structure can be determined from known circumstances, Content captures the predetermined text, while Logic defines the rules governing its inclusion, selection, calculation, and repetition. Document Automation applies this Content and Logic to produce a matter-specific document through predominantly deterministic behaviour, making both the text and the drafting method reusable.

AI Workflows make a method reusable where analysis must respond to the particular matter. Separating the work into defined tasks gives each analytical function the instructions, information, and output expectations it needs. Large Lan-

guage Models, AI models that process and generate language, apply that structure with probabilistic behaviour, allowing conclusions and wording to vary. The assets therefore complement each other: Document Templates preserve drafting choices settled in advance, while AI Workflows develop contextual analysis that can inform drafting or review the resulting documents.

Turning expertise into reusable assets requires a clear connection between the value sought, the result expected, and the method designed. Legal AI Design establishes that connection through three phases: Selection determines which work warrants investment; Specification defines what the assets must accomplish and how they will be evaluated; Design specifies how they should work. Together, the phases give development a business rationale, a defined legal purpose, and an implementation-ready design.

These three phases also establish the basis for evaluating the resulting assets after implementation. Legal AI Design organizes this relationship using the V-model, a development framework that pairs definition and design with corresponding evaluation. The descending side of the V moves from Opportunities through Requirements to Architecture and Design, with implementation at its base. The ascending side evaluates the implemented asset against those earlier decisions: Acceptance Criteria establish whether the Requirements are met, while Business Validation evaluates the business value realized through use. This connection gives each design decision a purpose and each expected outcome a basis for evaluation, linking the finished asset to the reasons for developing it.

Phase	Definition	Validation
Selection	Opportunities (Chapter 5)	Business Validation (Chapter 6)
Specification	Requirements (Chapter 7)	Acceptance Criteria (Chapter 8)
Design	Architecture and Design (Chapters 9–12)	

Selection

Selection directs productization toward legal work with the greatest potential for reuse. Opportunity Identification establishes where Document Templates and AI Workflows could capture recurring expertise and which Opportunities should proceed to Specification. Business Validation completes this connection on the validation side of the V-model: after implementation, it assesses whether the potential identified during Selection has translated into business value.

Opportunity Identification

Opportunity Identification reveals where capturing legal expertise once can benefit many matters. A Legal Domain is a practice area or legal niche; a Process is a recurring sequence or category of legal work within that domain. Reviewing the firm’s practice areas and representative matters identifies these Legal Domains and Processes, along with the documents or analytical tasks they repeatedly involve. The more frequently that work recurs, the greater the potential to reuse the captured expertise and the higher its priority for productization. Expected value can also influence an Opportunity’s attractiveness, but is excluded from this prioritization because

the value of individual Document Templates and AI Workflows typically varies less significantly than the frequency with which they can be reused.

Each Opportunity identifies a Document Template or AI Workflow and states its Objective: “Generate [document type] for [legal context]” for a Document Template, or “[verb] [noun] for [legal context]” for an AI Workflow. The inventory groups Opportunities by Legal Domain and Process, with both kinds of asset listed together in descending order of frequency within each Process. Their position expresses relative priority, giving the firm a practical basis for selecting the work that proceeds to Specification.

Opportunities

[legal domain]

[process]

[Document Template or AI Workflow] — [objective]

[...]

[...]

[...]

Business Validation

Business Validation establishes whether reusable legal expertise delivers business value in practice. Legal Assets can contribute to revenue-generating services or reduce the effort required for internal or non-billable work. Their value therefore depends on the revenue attributable to their use or the cost reduction achieved compared with equivalent manual work. Assessing these results after implementation connects

the potential identified during Selection to the business value actually realized.

Value per Use records the economic value realized each time a Legal Asset is used, based on the attributable revenue or resulting cost reduction. Annual Frequency records how often the asset is used during the year. Together, these measures establish its Annual Value: $\text{Annual Value} = \text{Value per Use} \times \text{Annual Frequency}$. The Business Validation record brings them together for each implemented asset, providing a common basis for evaluating the business value generated through repeated use of the captured expertise.

Business Validation

[Document Template or AI Workflow] — [objective]

Value per Use

[amount]

Annual Frequency

[uses per year]

Annual Value

[amount]

Specification

Once the Opportunities have been selected from the ordered inventory, Specification defines what each asset must accomplish. Requirements Determination describes the legal task, necessary inputs, and expected output, while Acceptance Criteria Setting establishes how satisfactory performance will be assessed. Together, these provide the targets that guide

Design and the standards against which the resulting implementation will be evaluated, connecting the required legal results to the business value sought.

Requirements Determination

Requirements Determination ensures that the asset is designed around a clearly defined legal need. It translates the selected Objective into the Legal Task, the circumstances in which it will be performed, and the Input Scope needed to complete it. Defining the matter information, legal knowledge, and other necessary inputs makes clear what work the asset must support and what information it requires.

Output Characteristics establish what a satisfactory result must contain or demonstrate. They give Design a clear target while leaving the mechanism for producing the result open. The Requirements record brings these characteristics together with the Legal Task and Input Scope, so the intended legal outcome can guide design decisions and provide the basis for Acceptance Criteria.

Requirements

[Document Template or AI Workflow] — [objective]

Legal Task

[description]

Input Scope

[description]

Output Characteristics

[description]

Acceptance Criteria Setting

Acceptance Criteria Setting allows quality to be assessed against agreed expectations. Requirements describe the intended result; criteria establish the evidence needed to accept it. Each criterion pairs Test Conditions with an Expected Result, specifying the representative inputs and legal circumstances under which the asset is assessed and what the evaluator must observe. This gives the firm a practical basis for deciding whether the implementation fulfils its purpose.

The criteria reflect how each asset produces its results. For Document Templates, they can address content, drafting logic and formatting. For AI Workflows, they define the substantive qualities an acceptable result must demonstrate while allowing variation in wording. Both types of asset can therefore be assessed against their Requirements using standards suited to their behaviour.

Acceptance Criteria

[Document Template or AI Workflow] — [objective]

[acceptance criterion]

Test Conditions

[description]

Expected Result

[description]

[...]

Document Automation Design

Once the Requirements and Acceptance Criteria are defined, Design determines how the assets will produce those results. For Document Templates, this begins with an Architecture that identifies shared Snippets (reusable components containing Content, Logic, and formatting) and Questions, which request the user input needed to generate the documents. Fields identify variable information to be inserted into the documents. Detailed Document Template Design and Snippet Design then specify the Content, Logic, formatting, Fields, and Questions needed to turn the firm's drafting expertise into instructions for implementation.

Document Template Architecture

Document Template Architecture allows shared drafting decisions to be developed and maintained consistently across documents. Comparing the selected templates' Requirements reveals recurring Content, Logic, and formatting that can become reusable Snippets. Questions capture the user input needed by the templates and their Snippets. Giving these components consistent names allows the same drafting and information needs to be addressed wherever they recur.

Mapping each Snippet and Question to the assets that use it makes their dependencies explicit before detailed design begins. This shows where a component must serve several documents and which assets a change will affect. The Architecture provides a common structure for Document Template Design and Snippet Design, extending the value of each drafting decision across the portfolio.

[snippet]

[template objective]

[...]

[question]

[template objective]

[...]

[snippet]

[...]

Document Template and Snippet Design

Document Template Design and Snippet Design preserve the lawyer's drafting choices as they are translated into automation. Each Document Template and Snippet has a Document: the source file containing the Content, formatting, and Logic to which automation will be applied. It keeps the wording and structure connected to the conditions, selections, calculations, and repetitions governing the text, giving the implementer the instructions needed to reproduce the intended drafting method.

Fields identify variable information drawn from the Client Case, while Questions connect required user input to its use in the Document. Document Templates also identify the Snippets they incorporate, with the Document specifying where these are inserted and when they apply. Snippets use the same design elements. Each deliverable brings the Document and its associated components together in a record accompanying the source file.

Document Template — [objective]

Document

[Content, formatting and Logic]

Fields

[field]

[...]

Questions

[question]

[...]

Snippets

[snippet]

[...]

[snippet]

Document

[Content, formatting and Logic]

Fields

[field]

[...]

Questions

[question]

[...]

Legal Intelligence Design

The same progression from Requirements to Architecture and detailed Design applies to AI Workflows, where the reusable asset captures an analytical method. AI Workflow Architecture identifies the Roles, which define reusable analytical approaches, and Instructions, which specify reusable assignments. AI Workflow Design then organizes them into connected Steps, individual units of Workflow execution that each apply a Role and an Instruction. Role Design and Instruction Design specify how each Step should approach its task and produce the required result.

AI Workflow Architecture

AI Workflow Architecture makes analytical approaches and assignments reusable across Workflows. Roles capture the approach associated with a Legal Function, the professional activity performed by the model, while Instructions describe the Legal Assignment, the legal work to be performed. Separating them allows the same approach to support different tasks and the same assignment to be reused wherever it is needed.

Mapping the Roles and Instructions to the Workflows that use them shows which components should be developed together. AI Workflow Design assembles those components into Steps, while Role Design and Instruction Design define their operation. This shared structure lets improvements to an analytical approach or assignment benefit every Workflow that uses it, building a common body of reusable expertise.

[role]

[workflow objective]

[...]

[...]

[instruction]

[workflow objective]

[...]

[...]

AI Workflow Design

AI Workflow Design makes complex analytical work manageable by separating its tasks and connecting their results. Numbered Steps distinguish changes in assignment, analytical function, or required output, with each Step receiving a Role and Instruction from the Architecture. This gives each part of the work a defined purpose and shows how the parts work together to address the Workflow's Requirements.

Knowledge Cases identify the reusable legal knowledge each Step needs to perform its task, such as legislation, case law, precedents, or internal guidance. Recording them at Step level allows the same Role and Instruction to be used with different knowledge sources. The Workflow record brings the Role, Instruction, and Knowledge Cases together for each Step under the overall Objective, making explicit how the method and its supporting knowledge contribute to the intended result.

AI Workflow — [objective]

Step 1

[role]

[instruction]

[knowledge cases or none]

Step 2

[role]

[instruction]

[knowledge cases or none]

[...]

Role and Instruction Design

Role Design makes the intended analytical approach explicit so it can be applied consistently across assignments. The Legal Function defines the professional activity the model performs, the Reasoning Perspective specifies how it should interpret and evaluate information, and Behavioural Constraints establish the boundaries of that work. Together, these three elements give each Step using the Role a common analytical basis.

Instruction Design defines what each assignment must accomplish and how its result should be expressed. The Legal Assignment describes the legal work to be performed, Applicable Rules guide its performance, and Required Output specifies the result to be produced. Supporting Knowledge Cases are identified separately for each Workflow Step, while matter information is supplied during execution. Separate

Role and Instruction records make both the analytical approach and the assignment reusable.

Role Design

[role]

Legal Function

[description]

Reasoning Perspective

[description]

Behavioural Constraints

[description]

Instruction Design

[instruction]

Legal Assignment

[description]

Applicable Rules

[description]

Required Output

[description]

Commercial and Strategic Advantage

Productization changes the economics of legal service delivery by shifting recurring effort into the design and maintenance of reusable assets. Each new matter benefits from

decisions and methods developed in advance, allowing firms to handle more matters with more consistent and predictable delivery. This also creates scope for pricing services according to the value delivered, while clients gain clearer expectations about the work, its timing, and its price.

As the portfolio develops, expertise becomes an institutional resource that can be applied across teams and improved over time. Shared components extend the value of each design effort, allowing new assets to build on methods and drafting decisions already captured. The firm's accumulated knowledge thus contributes to its delivery capacity beyond the individual matters in which that knowledge was first developed.

The strategic advantage lies in the portfolio itself. Competitors can acquire similar technologies, but the drafting decisions and analytical methods captured in the assets are developed through practice. They reflect the firm's legal knowledge, client experience, and working methods. Systematically capturing and improving that expertise gives the firm a distinctive resource for delivering its services. Legal AI Design provides the methodology for building it.

Contents

Introduction	1
I Foundations	10
1 Legal Work	11
2 Document Templates	23
3 AI Workflows	31
4 Design Methodology	40
II Selection	47
5 Opportunities	48
6 Business Validation	57
III Specification	65
7 Requirements	66
8 Acceptance Criteria	75

IV Document Automation Design	85
9 Document Template Architecture	86
10 Document Template Design	95
V Legal Intelligence Design	105
11 AI Workflow Architecture	106
12 AI Workflow Design	115
Conclusion	127
Glossary	135

Introduction

Legal expertise is valuable because it enables lawyers to recognize what matters, determine what should be done, and produce legal work that responds to a client's particular circumstances. Yet the way legal services are traditionally organized makes that expertise difficult to reuse independently of the lawyers who possess it. A lawyer applies expertise to one matter, produces the required work, and then applies much of the same expertise again when a similar matter arrives. Legal AI Design starts from a different premise: legal expertise can be systematically transformed into reusable legal assets that allow it to be applied repeatedly across matters. This creates the foundation for productized legal services.

This book develops a methodology for making that transformation possible. It focuses on two reusable legal assets: Document Templates, which support Document Automation, and AI Workflows, which support Legal Intelligence. It then provides a structured process for designing those assets, progressing from Selection through Specification to Design. Together, these elements provide a practical framework for moving legal expertise from the individual lawyer into systems through which that expertise can be reused.

The Business Model Problem

The traditional legal services business model is built around a simple relationship between expertise and time. Clients

engage lawyers because those lawyers possess legal knowledge, professional judgment, and experience that the clients themselves do not possess. The lawyer applies that expertise to the client's particular matter, and the client pays for the professional work required to do so.

This model has proved remarkably durable. It is also inherently difficult to scale. Legal expertise resides primarily in individual lawyers and is repeatedly applied through their work on individual client matters. A lawyer may have drafted the same type of agreement dozens of times, analyzed the same category of legal problem repeatedly, or handled many disputes that follow broadly similar patterns. Each new matter may nevertheless require that lawyer, or another lawyer with comparable expertise, to perform much of that legal work again.

The consequence is that the capacity of a traditional legal practice remains closely connected to the time available from its lawyers. A lawyer can become more experienced and more efficient, and a law firm can employ more lawyers, improve its internal processes, and use better technology. But the underlying commercial model still largely depends on selling the application of expertise through lawyer time. Hourly billing is the clearest expression of this relationship: the unit through which legal expertise is commercialized is the professional hour.

This creates an important distinction between expertise itself and the way expertise is delivered. Legal expertise may have been developed over years or decades, but its commercial application is often recreated one matter at a time. The expertise therefore has substantial intellectual value without necessarily becoming an asset that can be deployed independently of the lawyer who possesses it.

Productization changes this relationship.

A productized legal service does not require every element of the service to be reconstructed for every new matter. Instead, recurring expertise is captured in reusable legal assets that can be applied repeatedly. The lawyer's expertise remains fundamental, but part of that expertise is transferred from individual professional performance into designed systems.

Consider a recurring legal document. In a traditional model, a lawyer may begin with a previous document, modify it for a new matter, identify the relevant alternatives, insert client information, adapt clauses, check internal consistency, and review the completed document. Much of the lawyer's expertise is embedded in those decisions. If that expertise can instead be captured in a sufficiently sophisticated Document Template, the template becomes a reusable legal asset through which the same expertise can be applied across many matters.

The same principle applies to analytical legal work. A lawyer reviewing a dispute may repeatedly identify relevant facts, organize them around recurring legal issues, identify missing evidence, retrieve applicable legal rules, apply those rules, assess risks, and prepare advice. Generative AI can perform parts of this work, but merely giving a lawyer access to a general-purpose AI system does not capture the lawyer's methodology. If the relevant legal reasoning process is instead decomposed into a controlled sequence of instructions, roles, information sources, and outputs, it can become an AI Workflow that supports repeated application of that expertise.

Productization, in this sense, does not mean turning every legal service into a rigid commodity. Nor does it require removing lawyers from legal work. It means identifying ex-

pertise that is sufficiently recurring to be captured and systematically transforming that expertise into reusable legal assets.

Those assets change the economics of legal service delivery. Once created, they can support repeated application of expertise without requiring the underlying method to be reconstructed from the beginning each time. This can make services more standardized, increase consistency, reduce the amount of repetitive professional work, and create greater predictability in how the service is delivered. It can also support pricing models in which price is linked more closely to the service delivered than to the number of professional hours required to produce it.

The central business model problem addressed by this book is therefore not simply that lawyers need better technology. It is that legal expertise remains insufficiently productized. Technology becomes strategically important when it provides a means of transforming that expertise into reusable assets.

The Reusable Legal Assets

Legal AI Design focuses on two principal types of reusable legal assets: Document Templates and AI Workflows. They address different forms of legal work and rely on different technological capabilities, but they serve the same broader purpose. Both capture aspects of legal expertise so that they can be applied systematically across matters.

Document Templates support Document Automation. They are appropriate where legal work produces documents whose structure and content recur across matters and where relevant variations can be identified in advance. Agreements, letters, pleadings, reports, notices, corporate documents, and

many other legal outputs frequently contain a combination of recurring text, variable information, optional provisions, conditional clauses, repeated data, and calculated elements.

A properly designed Document Template captures this structure. Instead of treating every document as a new drafting exercise, the lawyer identifies which parts are stable, which depend on matter-specific information, which alternatives are available, and under what conditions particular content should appear. Document Automation can then assemble a matter-specific document from that reusable structure.

The value of the template lies not merely in reusing text. A sophisticated Document Template captures drafting decisions. It represents the lawyer's understanding of the document: what information matters, which provisions belong together, which alternatives exist, when those alternatives should be used, and how changes in one part of the document affect another. The template therefore becomes a structured representation of legal drafting expertise.

Not all legal work can be represented effectively in this way. Many legal tasks depend on unstructured facts, variable documents, legal research, interpretation, comparison, analysis, and judgment. The sequence of work may be recognizable even where the content cannot be predetermined. These forms of work require a different reusable asset.

AI Workflows support Legal Intelligence. An AI Workflow structures a recurring form of knowledge work as a sequence of defined steps. Each step performs a particular function, such as extracting facts, identifying gaps, identifying legal issues, retrieving legal rules, applying those rules, identifying risks, drafting an output, or reviewing the reasoning that produced it. Different roles and instructions can be assigned to different steps, and the output of one step can become the

input for another.

The objective is not simply to write a better prompt. A prompt typically asks an AI system to perform a task at a particular moment. An AI Workflow captures a reusable method for performing a category of legal work. It makes the sequence of that work explicit and separates complex legal tasks into components that can be defined, controlled, reviewed, and improved.

Document Templates and AI Workflows therefore capture different dimensions of legal expertise. Document Templates are particularly suited to legal work where the desired output can be modeled through predefined structure, variables, conditions, and reusable drafting. AI Workflows are particularly suited to legal work where the method is recurring but the information being processed and the conclusions to be developed vary substantially from matter to matter.

The distinction is important because legal technology should follow the structure of the legal work. Attempting to solve every problem with generative AI ignores the value of deterministic automation where the required logic can be specified in advance. Conversely, attempting to reduce complex analytical work to fixed templates ignores the flexibility required when facts, documents, legal questions, and reasoning vary between matters.

Legal AI Design therefore does not begin with a preferred technology. It begins with legal work. The purpose is to determine what expertise is being applied, what parts of that expertise can be made reusable, and what type of legal asset can represent it effectively.

The Design Methodology

Recognizing that legal expertise can be transformed into reusable assets does not explain how to perform that transformation. Lawyers possess expertise in forms that are often implicit. Experienced practitioners may know instinctively which questions to ask, which facts deserve attention, which clauses require modification, which risks matter, or which analytical steps should follow one another. That expertise must be made explicit before it can be embedded reliably in a reusable legal asset.

Legal AI Design provides a structured methodology for doing so. The methodology progresses through three stages: Selection, Specification, and Design. Each stage addresses a different design problem.

Selection determines what should be created. Not every legal task should be automated, and not every recurring activity justifies investment in a reusable asset. The first challenge is therefore to identify suitable legal work, understand how it recurs, determine where productization can create value, and select the appropriate type of asset. This prevents technology from becoming the starting point and focuses investment on legal work for which reuse has a meaningful purpose.

Specification determines what the selected asset must accomplish and how this will be evaluated. Once a productization opportunity has been selected, the underlying legal expertise must be examined systematically. Requirements define the expected behavior and characteristics of the legal asset, while corresponding Acceptance Criteria establish how it will be determined whether those Requirements have been satisfied. Together, they translate professional practice into an explicit specification without yet determining how the asset will be implemented.

Design determines how the selected legal assets should work. It moves from their overall Architecture—the structure and relationships between the selected assets—to the detailed Design of each individual Document Template or AI Workflow. At this stage, the Requirements are translated into specifications sufficiently detailed to guide implementation.

The progression from Selection to Specification to Design is deliberate. Beginning directly with technology encourages premature solutions: automating a document before understanding its variations, constructing an AI prompt before decomposing the underlying legal reasoning, or investing in a workflow before establishing whether the work is sufficiently recurring to justify productization. Legal AI Design instead separates the decision about what to build from the definition of what it must accomplish and from the design of how it will operate.

This makes productization a design discipline rather than an ad hoc technology exercise. The reusable legal asset is not merely software configured by a technologist. It is a representation of professional expertise that must be deliberately extracted, structured, and encoded. Lawyers therefore have a central role in its design because the quality of the resulting system depends on the quality with which the underlying legal expertise has been understood and transformed.

The methodology developed throughout this book provides the structure for that transformation. It connects the economics of productized legal services with the practical design of the assets through which productization becomes possible.

Conclusion

Legal expertise does not become scalable merely because a law firm adopts automation or generative AI. It becomes

more reusable when recurring professional knowledge and methods are deliberately transformed into legal assets that can apply that expertise repeatedly. Document Templates and AI Workflows provide two mechanisms for doing so; Legal AI Design provides the methodology for designing them.

Before that methodology can be applied, however, its foundations must be understood. We first need to examine the nature of the legal work being transformed and the technologies through which reusable legal assets can operate. Those foundations provide the vocabulary and conceptual framework required for the Selection, Specification, and Design stages that follow.

Part I

Foundations

1 Legal Work

Legal expertise traditionally resides primarily in individual lawyers, who repeatedly apply their knowledge, judgment, and methods to client matters. Even when similar work recurs, its performance requires lawyer time. Traditional legal services largely commercialize that expertise through time, typically billed by the hour, tying the firm's capacity and revenue closely to the professional hours available. Productization addresses this limitation by transforming legal expertise into reusable Legal Assets: Document Templates and AI Workflows. Legal AI Design provides the methodology for this transformation.

The two kinds of Legal Asset capture different aspects of expertise. Document Templates preserve drafting decisions in predefined Content and Logic. AI Workflows preserve analytical methods in structured Steps that respond to the information in each matter. Both allow expertise to be applied repeatedly with less recurring lawyer effort. Understanding how legal work is organized and how these assets complement each other provides the foundation for designing them.

A Legal AI System supports this work through three complementary capabilities: Case Management, Document Automation, and Legal Intelligence. Case Management provides the operational environment in which legal work takes place. Document Automation applies Document Templates to generate matter-specific documents. Legal Intelligence supports

contextual analysis and language-based legal work, including the execution of AI Workflows and assistance with the development of reusable assets.

Combining these capabilities connects matter information, reusable legal knowledge, drafting decisions, and analytical methods. The Legal AI System provides the environment in which expertise can be captured, applied to individual matters, and refined for further use.

Case Management

Legal work is performed in the context of matters. A lawyer advises a client, negotiates a transaction, pursues a claim, defends proceedings, conducts an investigation, or performs another legal service in relation to a particular set of circumstances. That context must be organized before technology can effectively support the legal work performed within it.

Case Management provides this operational foundation. At its core are Client Cases: structured environments containing the information and work associated with a particular client matter. A Client Case can contain the parties involved, documents received or created, correspondence and other communications, tasks, deadlines, time registrations, costs, and the work product produced during the matter. Rather than treating these elements as disconnected records spread across different applications and locations, the Client Case brings them together around the legal matter to which they belong.

This organization is important because legal information derives much of its meaning from context. A document is not simply a file. It may be a contract governing the relationship between particular parties, evidence supporting a particular claim, correspondence affecting a procedural deadline, or

an earlier version of a document currently being negotiated. Similarly, an email, task, deadline, time entry, or cost becomes meaningful because of its relationship to a particular matter.

Client Cases therefore provide more than administrative convenience. They establish the information environment within which legal work can be performed. When a lawyer drafts a document, reviews correspondence, analyzes evidence, or uses artificial intelligence to answer a question about a matter, the relevant information can be obtained from the Client Case rather than assembled repeatedly from separate sources.

Not all legal information, however, belongs exclusively to one client matter. Lawyers also rely on legal knowledge that can be reused across many matters. A statute may govern hundreds of cases. A judgment may establish a principle relevant to a category of disputes. Legal doctrine may provide an analytical framework that lawyers repeatedly apply to different factual situations.

A Legal AI System can organize this reusable knowledge through Knowledge Cases. Unlike Client Cases, which represent individual client matters, Knowledge Cases contain collections of legal knowledge intended for use across matters. They can include legislation, case law, legal doctrine, internal legal analyses, guidance, and other relevant sources of knowledge.

The distinction is therefore functional. Client Cases organize the information and work associated with a specific matter. Knowledge Cases organize legal knowledge that can be reused across Client Cases. Together, they provide the information foundation on which Document Automation and Legal Intelligence can operate.

Document Automation

A substantial part of legal work consists of producing documents. Agreements, letters, pleadings, opinions, reports, corporate documents, and many other forms of legal work product contain elements that recur across matters. Although the factual circumstances change, portions of the resulting document can often be determined in advance.

Document Automation makes this repeatability explicit. It applies Document Templates to information from Client Cases to generate matter-specific documents. A precedent preserves earlier wording; a Document Template also captures the decisions governing how that wording changes to suit the matter. Content captures the predetermined text, while Logic defines the rules governing its inclusion, selection, calculation, and repetition. Applying this Content and Logic produces predominantly deterministic behaviour: the same relevant inputs and rules yield the same drafting choices.

Consider a relatively simple legal letter. The identities and addresses of the parties may already be stored in the Client Case. The appropriate introductory text may depend on the type of matter. A paragraph may be required only if a particular condition applies. Dates, amounts, references, and other information may be calculated or inserted automatically. Where these relationships can be specified beforehand, there is little value in requiring a lawyer to reproduce the same drafting process manually for every matter.

Document Automation therefore performs the parts of drafting that can be determined through predefined Content and Logic. Once those rules have been designed and incorporated into a Document Template, they can be applied repeatedly. This can make recurring drafting work faster while improving consistency and reducing the risk that standard content

is omitted, copied incorrectly, or applied inconsistently.

The boundary of Document Automation is equally important. Not every part of a legal document can or should be determined in advance. Some drafting decisions depend on facts whose legal significance cannot be reduced to predefined conditions. Others require interpretation, argumentation, judgment, or the weighing of competing considerations. A lawyer drafting a procedural submission, for example, may be able to automate the identification of the parties, procedural information, standard headings, and recurring formal language while still needing to develop an argument specifically adapted to the evidence and legal issues in the case.

Legal AI Design therefore distinguishes between drafting choices settled in advance through Content and Logic and analysis that must respond to the particular matter. Document Automation is particularly effective for the first category. Attempting to force the second category into increasingly complex deterministic templates can make automation difficult to design, maintain, and use.

Document Templates can capture the predetermined parts of a document while leaving contextual analysis to a lawyer or an AI Workflow. For example, a template can supply the structure and standard provisions of a procedural submission, while a Workflow develops an analysis of the evidence that informs its argument. The assets can therefore contribute to the same work product.

Detailed design makes the captured drafting method ready for implementation. Fields identify variable information drawn from the Client Case, while Questions request the user input needed to generate the document. Snippets are reusable components containing Content, Logic, and formatting.

Document Template Architecture identifies shared Snippets and Questions by comparing the selected templates' Requirements. It gives these components consistent names and maps each Snippet to the Document Templates that use it, and each Question to the Document Templates and Snippets that need its input. These dependencies show where a drafting decision or information need is shared and which assets a change will affect.

Document Template Design records the Document, Fields, Questions, and incorporated Snippets; Snippet Design records the Document, Fields, and Questions for each Snippet. The Document is the source file containing the Content, formatting, and Logic to which automation will be applied. For a Document Template, it also specifies where Snippets are inserted and when they apply. These design records accompany the source files so that the intended drafting method can be implemented consistently.

Legal Intelligence

Much of legal work cannot be fully expressed as predefined Content and Logic. Lawyers interpret language, identify relevant information, compare sources, analyze relationships between facts and legal rules, develop arguments, assess documents, and communicate conclusions. The appropriate output often depends on the meaning of the information available rather than on a predetermined set of conditions.

Legal Intelligence supports this language-based legal work. It can assist with activities such as legal research, analysis, drafting, reviewing, summarizing, comparing, and extracting information. The intended qualities of the result can be specified in advance, even though the conclusions and wording must respond to the particular matter. This distinction allows contextual work to have clear Requirements without

predetermining every part of its output.

The information available to Legal Intelligence is consequently important. When supporting work on a particular matter, it can use case-specific information from the Client Case. A lawyer might ask it to summarize a collection of correspondence, compare different versions of an agreement, identify relevant provisions in documents, extract facts from a case file, review a draft, or assist in developing an argument based on the information available in the matter.

Where the work also requires reusable legal knowledge, Legal Intelligence can use Knowledge Cases. The case-specific information contained in a Client Case can then be considered together with legislation, case law, legal doctrine, or other knowledge maintained independently of that particular matter. The system can thereby combine two forms of context: information about the matter and legal knowledge that may be applicable to it.

Large Language Models, AI models that process and generate language, support this work with probabilistic behaviour. Their conclusions and wording can vary even when the same task and information are supplied. Defined instructions and relevant sources guide the analysis, while Acceptance Criteria establish the substantive qualities an acceptable result must demonstrate. Lawyers review the resulting work and determine whether it is legally and factually appropriate.

AI Workflows make the analytical method reusable by separating the work into connected, numbered Steps, each an individual unit of Workflow execution. Changes in assignment, analytical function, or required output distinguish the Steps. Each Step receives a Role defining its analytical approach, an Instruction defining its assignment, and the Knowledge Cases it needs, or an explicit indication that none

are required. Matter information is supplied during execution. For example, a Workflow might extract relevant facts from correspondence, assess those facts against applicable legal rules, and then produce an analysis to inform drafting or review a resulting document. Connecting these tasks and their results gives each analytical function the instructions, information, and output expectations it needs to contribute to the overall result.

AI Workflow Architecture identifies reusable Roles and Instructions and maps them to the Workflows that use them. Separating the analytical approach from the assignment allows a Role to support different tasks and an Instruction to be reused wherever its assignment is needed. AI Workflow Design assembles these components into Steps and records the Role, Instruction, and Knowledge Cases for each Step under the Workflow's overall Objective.

Role Design specifies the Legal Function, the professional activity performed by the model; the Reasoning Perspective, how it should interpret and evaluate information; and Behavioural Constraints, the boundaries of its work. Instruction Design specifies the Legal Assignment, the legal work to be performed; Applicable Rules, which guide its performance; and Required Output, the result to be produced. Separate records capture each Role and Instruction. Keeping supporting Knowledge Cases at Step level allows the same Role and Instruction to be reused with different knowledge sources.

Legal Intelligence can also contribute to the creation of Legal Assets themselves. Developing Document Templates requires lawyers to identify recurring language, alternative clauses, relevant conditions, variables, and relationships between different parts of a document. Legal Intelligence can assist by analyzing existing documents, proposing reusable formulations, identifying variations, drafting clauses, and

helping translate Requirements into Content and Logic. For AI Workflows, it can help separate an analytical method into Steps and formulate Roles and Instructions for review by the lawyer.

The same applies when assets must be maintained. Changes in legislation, case law, organizational practice, or drafting standards may require Content and Logic, analytical methods, or supporting Knowledge Cases to be reconsidered. Legal Intelligence can assist lawyers in identifying and preparing those changes. The resulting assets must be evaluated against their Requirements and Acceptance Criteria before the revised expertise is relied on across matters.

Legal Intelligence therefore has two related roles. It can support lawyers in creating and maintaining reusable legal expertise, and it can support contextual legal work within individual matters. Both roles become more valuable when Legal Intelligence operates within the same information environment as Case Management and Document Automation.

The Legal AI System

Case Management, Document Automation, and Legal Intelligence become most useful when they are understood as complementary capabilities within one Legal AI System. Each addresses a different requirement of legal work, and each provides inputs that can be used by the others.

Case Management provides the operational environment. A Client Case contains the parties, documents, communications, tasks, deadlines, time, costs, and work product associated with a matter. Knowledge Cases provide reusable collections of legislation, case law, legal doctrine, and other knowledge that can support work across multiple Client Cases.

Reusable legal expertise can then be developed within this

environment. Lawyers identify recurring drafting decisions for Document Templates and analytical methods for AI Workflows. Legal Intelligence can assist in developing and maintaining both kinds of Legal Asset. Shared Snippets and Questions extend reuse across Document Templates, while shared Roles and Instructions extend reuse across AI Workflows. Lawyers define the intended legal results and review the assets against the agreed standards.

When a Document Template is used in a particular Client Case, Document Automation combines that reusable asset with information from the matter. It retrieves or receives the relevant case information and applies the template's predefined Content and Logic. The result is a document in which the parts that can be determined in advance have already been generated.

AI Workflows apply reusable analytical methods to the same matter information, drawing on the Knowledge Cases identified for each Step where appropriate. A Workflow can analyze facts before a document is generated, develop contextual text to inform drafting, or review a document afterwards. Legal Intelligence can also assist lawyers with individual tasks outside a predefined Workflow. The lawyer reviews and refines the resulting work in its legal and factual context.

The completed work product becomes part of the Client Case, where it can serve as an input to subsequent work. Reusable legal knowledge is organized in Knowledge Cases, while Document Templates and AI Workflows are retained as Legal Assets within the system. Retaining both work and assets allows experience from individual matters to inform the refinement of expertise used across matters. Retaining a matter document preserves its content; capturing its drafting decisions or analytical methods in a Legal Asset makes the underlying expertise reusable as a method.

This creates a coherent model of legal work. Case Management determines where matter-specific information and reusable legal knowledge are organized. Document Automation applies the predefined Content and Logic captured in Document Templates. Legal Intelligence supports contextual legal work, including the analytical methods captured in AI Workflows, and can assist in developing both kinds of asset.

This division of work gives each capability a clear purpose. Predetermined drafting choices can be applied consistently through Document Automation, while Legal Intelligence addresses analysis whose conclusions depend on meaning and context. Capturing that analysis in an AI Workflow makes the method reusable without making its behaviour deterministic.

Conclusion

Case Management, Document Automation, and Legal Intelligence provide the operational foundation for productizing legal expertise. Document Templates and AI Workflows turn drafting decisions and analytical methods into reusable Legal Assets. The Legal AI System provides the environment in which these assets can be developed, applied to individual matters, and refined for further use.

Repeated use can reduce recurring lawyer effort, support more consistent delivery, and create scope for more predictable pricing where the resulting services are offered to clients. As the portfolio develops, the firm retains expertise as an institutional resource that can be maintained, improved, and applied across matters and teams. Shared components allow new assets to build on drafting decisions and methods already captured. The resulting portfolio reflects the firm's legal knowledge, client experience, and working methods, giving it a distinctive resource for delivering legal

services.

The next chapter examines Document Templates, the first of the two kinds of Legal Asset designed through the Legal AI Design methodology. It considers how recurring legal drafting can be transformed from work repeatedly performed in individual matters into structured expertise that can be maintained, improved, and systematically reused.

2 Document Templates

The previous chapter introduced Document Automation as one of the capabilities of a Legal AI System. Document Automation generates documents by applying the reusable drafting expertise captured in Document Templates to information from a Client Case. The capability performs the generation; the Document Template preserves the expertise that can be applied again across matters.

A Document Template is a reusable Legal Asset that captures established drafting decisions in predefined Content and Logic. Content captures the predetermined text, while Logic defines the rules governing its inclusion, selection, calculation, and repetition. Formatting specifies how the resulting document should be presented. Fields identify variable information to be inserted, Questions request the user input needed for generation, and Snippets provide reusable drafting components.

The generated contract, letter, pleading, or report is a document for a particular Client Case and purpose. The Document Template captures the repeatable drafting method from which such documents can be produced. Making that method reusable reduces the need for lawyers to reconstruct the same drafting decisions each time similar work arises.

From Precedents to Document Templates

Lawyers have always reused documents. Most legal practices maintain collections of precedents, model documents, previous agreements, standard letters, pleadings, and other examples that lawyers can retrieve when similar work arises. These preserve valuable wording and examples of earlier work.

A lawyer using a precedent identifies a previous document that appears sufficiently similar to the document now required, creates a copy, and adapts it to the new circumstances. Names are replaced, dates are changed, provisions are added or removed, amounts are updated, and language is modified where necessary. The previous document provides a useful starting point, but much of the expertise required to adapt it remains with the lawyer.

That expertise includes knowing what information must be changed, which provisions are relevant, which alternatives should be selected, which clauses depend on particular circumstances, and whether changes in one part of the document require corresponding changes elsewhere. The precedent preserves the result of earlier drafting decisions without necessarily making those decisions explicit. A lawyer using it must therefore reconstruct at least part of the method that produced it.

A Document Template also captures the decisions governing how wording should change to suit the matter. Variable information is identified through Fields. Content that depends on known circumstances is governed by Logic. Repeated structures can be generated systematically, while common components can be maintained as Snippets and reused across templates. Questions identify the user input needed to apply these drafting choices.

Existing documents remain an important source for this work. Comparing precedents can reveal common wording, alternative provisions, recurring variables, and the circumstances governing their use. Legal AI Design turns those observations into explicit instructions for implementation. The resulting asset preserves both the text and the drafting method needed to produce appropriate documents repeatedly.

Content and Formatting

Content is the predetermined text captured in a Document Template or Snippet. It may include headings, standard provisions, explanatory text, definitions, and the wording of signature blocks. Establishing this text in advance gives generated documents a consistent drafting foundation.

Predetermined text does not have to appear in every generated document. Several alternative clauses can all be defined as Content, with Logic selecting the appropriate one. A provision can be drafted in advance but included only when specified circumstances apply. Content therefore describes the text available for use; Logic determines how it is used in the particular matter.

Formatting specifies presentation and layout alongside Content and Logic. It can include heading styles, paragraph numbering, spacing, tables, and signature arrangements. These choices help preserve the intended structure and readability when provisions are selected, omitted, or repeated. Separating the concepts makes clear that wording, presentation, and the rules governing variation each require attention during Design.

In Document Template Design, the Document is the source file containing the Content, formatting, and Logic to which

automation will be applied. It keeps the wording and structure connected to the instructions governing their use. This source Document forms part of the design deliverable and is distinct from the document later generated for a Client Case.

Fields and Questions

Fields identify variable information drawn from the Client Case for insertion into the Document. Typical examples include party names, addresses, dates, reference numbers, amounts, contact details, deadlines, and case identifiers. Rather than requiring a lawyer to locate and replace each value in a copied precedent, the design identifies where that information belongs.

Case Management supplies the matter information used in generation. Information about clients, counterparties, transactions, invoices, claims, or other records can therefore flow into a generated document from the Client Case. A party name used in several places can be drawn from the same matter information wherever it is needed, reducing repeated entry and the risk of inconsistent replacements.

Questions request the user input needed to generate the document. They connect that input to its use in the Document Template or Snippet. A Question might ask the user to supply information needed for the current document or to choose between drafting alternatives where the relevant choice is not already established by the Client Case information.

Fields and Questions have distinct purposes. A Field identifies variable information for insertion; a Question identifies input to be requested from a user. An answer may supply a needed value or determine which Content is selected by Logic. Questions can therefore affect both the information appearing in a document and the drafting choices applied to

it.

For example, a template may draw a party's name and address from the Client Case through Fields, while a Question asks whether a specified optional provision should apply. Logic then includes or omits that provision according to the answer. The same template combines existing matter information with the user input required for that generation.

Questions can also serve several Document Templates or their Snippets. Identifying the same information need consistently allows it to be addressed wherever it recurs. Document Template Architecture makes these dependencies explicit before detailed design begins, while the individual designs connect each Question to its use in the relevant Document.

Logic

Not every difference between legal documents can be handled by inserting different values. Documents frequently change because particular provisions, structures, calculations, or alternatives are appropriate only in particular circumstances. Logic captures the predefined rules governing these variations.

These rules can govern inclusion, selection, calculation, and repetition. A provision may be included only if a particular condition is satisfied. One of several predetermined alternatives may be selected according to the type of party or transaction. An amount may be calculated from underlying values. A section may repeat for every relevant person, invoice, asset, claim, or other record. Relationships between these rules can ensure that a choice in one part of the document is reflected elsewhere.

These mechanisms allow a single Document Template to represent multiple predictable variations of a document. Instead

of maintaining separate versions for every combination of circumstances, the relevant differences can be specified as Logic within the reusable asset.

This captures drafting expertise directly. A lawyer may know that a provision must be inserted whenever a defined circumstance exists, that a calculation must follow an established formula, or that different contractual alternatives apply to different categories of matter. Where that decision can be expressed as an explicit rule, it can become repeatable behaviour within the template. The lawyer's drafting decision is then preserved for consistent application whenever the specified circumstances occur.

Document Automation applies this Content and Logic through predominantly deterministic behaviour. Where the relevant information and rules are fixed, the resulting drafting choices are determined in advance. Extensive conditions, calculations, selections, and repetitions can represent a sophisticated drafting method while retaining this characteristic.

There is an important boundary to what these rules capture. The designer must be able to specify the appropriate result for defined inputs or circumstances. Where the work requires contextual interpretation that cannot adequately be expressed through predetermined alternatives or rules, an AI Workflow can provide analysis to inform drafting or review the resulting document. The two kinds of Legal Asset can support the same matter, each capturing the expertise suited to its behaviour.

Snippets

Some drafting decisions recur across several Document Templates. The same provision, explanation, signature arrange-

ment, or conditional structure may be required in many documents. Maintaining each occurrence independently duplicates work and makes consistency harder to preserve.

Snippets are reusable drafting components containing Content, Logic, and formatting. A Snippet may be a standard paragraph or a more extensive structure with conditional provisions, calculations, or repeated sections. Its defining characteristic is that it can be developed and maintained as a component for reuse in Document Templates.

If the same clause is required in ten templates, maintaining ten separate copies creates ten potential points of inconsistency. A legal change may require each copy to be identified and updated correctly. Capturing the clause in a shared Snippet allows the drafting decision to be maintained together with the rules and formatting governing its use. Document Template Architecture identifies the templates that depend on that component, making the effects of a change visible.

Reuse can extend beyond common wording. A Snippet may use Fields for party information, Questions for required user input, and Logic to select or repeat Content. A signature component, for example, can combine predefined text with the relevant party details and rules governing the signature arrangement.

The Document for each template specifies where its Snippets are inserted and when they apply. The Snippets' own detailed designs specify their Content, Logic, formatting, Fields, and Questions. This preserves a clear connection between the reusable component and the documents that use it.

Snippets extend reuse within and across Document Templates. They are components of the drafting assets addressed by the methodology; Document Templates and AI Workflows remain the two kinds of Legal Asset. Shared components

allow the expertise captured in one design effort to benefit several assets and make the portfolio easier to maintain consistently.

Conclusion

Document Templates capture established drafting decisions so they can be applied repeatedly across Client Cases. Content preserves predetermined text, Logic governs its use, and formatting defines its presentation. Fields identify variable matter information, Questions request required user input, and Snippets make shared drafting components reusable across documents.

The combination varies with the work. A simple template may rely mainly on predefined text and a few Fields. Another may contain extensive Logic or draw heavily on shared Snippets and Questions. In each case, the template makes the established drafting method explicit so it can be applied consistently to the circumstances of individual matters.

Each generated document serves an individual matter. The reusable value lies in the drafting decisions retained in the Document Template and its components. Applying those decisions across matters can reduce recurring lawyer effort, improve consistency, and support more standardized and scalable legal work. Where the resulting services are offered to clients, this also creates scope for more predictable pricing.

AI Workflows complement these assets by capturing analytical methods that respond to the particular matter. Their contextual analysis can inform drafting or review the documents produced through Document Automation. The next chapter examines how these methods become the second kind of reusable Legal Asset addressed by Legal AI Design.

3 AI Workflows

The previous chapter established how Document Templates capture legal expertise through predefined Content and Logic applied with predominantly deterministic behaviour. They are particularly effective where the structure of the required output is known in advance and the relevant variations can be expressed through variables, conditions, calculations, and reusable clauses. Much legal work, however, does not operate in this way.

Legal professionals also work extensively with language that cannot be predetermined. They summarize case files, extract relevant facts from documents, classify information, compare contractual provisions, analyze legal issues, reformulate arguments, and draft text in response to circumstances that vary from case to case. The required output depends not merely on predefined data points, but on the meaning of the information presented. For this type of work, legal expertise requires a different mechanism.

An AI Workflow is a reusable Legal Asset that captures an analytical method in structured, numbered Steps. Each Step applies a Role and an Instruction and identifies its supporting Knowledge Cases, or indicates that none are required. Matter information from the Client Case is supplied during execution, and human involvement can be incorporated where appropriate.

The purpose of an AI Workflow is therefore not simply to

provide access to artificial intelligence. It is to transform the general language capabilities of a Large Language Model into a structured process for performing a defined form of legal work. To understand how this transformation takes place, it is first necessary to understand what a Large Language Model contributes and, equally importantly, what it does not provide by itself.

Large Language Models

A Large Language Model, or LLM, is an AI model that processes and generates language, trained on very large quantities of text. At a practical level, it generates language by identifying statistical patterns in the information it has processed and predicting what language is appropriate in the context it receives. It does not retrieve a predetermined answer from a database or execute legal reasoning in the deterministic manner of a traditional software rule.

This probabilistic approach makes Large Language Models particularly useful for work in which language itself is the object of the task. An LLM can draft a description based on facts, summarize a lengthy document, extract particular information from unstructured text, classify documents or passages, compare different versions of a contract, or reformulate existing language for a different purpose. These capabilities correspond closely to many activities performed by legal professionals every day. Conclusions and wording can vary even when the same task and information are supplied. Acceptance Criteria therefore assess the substantive qualities an acceptable result must demonstrate while allowing variation in wording.

The significance of this capability should not be confused with autonomous legal expertise. A general-purpose Large Language Model can generate plausible legal language, but

plausibility is not the same as reliable legal work. The quality of the result depends heavily on how the model is used. In practice, three recurring problems must be addressed when applying Large Language Models to legal work: objective ambiguity, information deficiency, and context overload.

Objective ambiguity arises when the model does not have a sufficiently precise understanding of what it is expected to do. An instruction such as “analyze this contract” leaves many questions unanswered. Should the model identify obligations, assess enforceability, compare provisions against a precedent, identify missing protections, or draft proposed amendments? Different interpretations can produce materially different results.

Information deficiency arises when the model does not have the information required to perform the task. A model cannot reliably apply a particular legal rule if that rule is unavailable to it, nor can it assess a client’s position without the relevant facts and documents. General knowledge embedded in the model cannot substitute for the authoritative and case-specific information on which professional legal work depends.

Context overload creates the opposite problem. Giving the model every available document, precedent, communication, legal source, and piece of background information does not necessarily improve the result. Irrelevant information can obscure the information that matters. Effective legal AI therefore requires not only information, but the right information for the particular task being performed.

An AI Workflow addresses these three problems systematically. Roles and Instructions define the analytical approach and assignment. Knowledge Cases and Client Cases provide the information. Retrieval-Augmented Generation can help

provide focused context.

Roles and Instructions

Roles and Instructions address objective ambiguity by making explicit both the function the model is performing and the task it must execute. Instead of asking a general-purpose model to “work on” a legal matter, an AI Workflow defines what the model is doing at each stage.

A Role defines a reusable analytical approach through three elements: Legal Function, the professional activity the model performs; Reasoning Perspective, how it should interpret and evaluate information; and Behavioural Constraints, the boundaries of its work. Depending on the workflow, the model might operate as a fact analyst, legal researcher, legal analyst, risk analyst, legal drafter, or legal reviewer. The Role establishes the functional perspective from which the task is approached and helps constrain the model to the type of work expected at that particular point in the process.

An Instruction defines a reusable assignment through three elements: Legal Assignment, the legal work to be performed; Applicable Rules, which guide its performance; and Required Output, the result to be produced. Supporting Knowledge Cases are identified separately for each Step, while matter information is supplied during execution. This allows the same Role and Instruction to be used with different knowledge sources. An instruction to extract facts, for example, can define which factual categories must be used, require preservation of conflicting information, prohibit the inference of missing facts, and prescribe a structured output identifying gaps and supporting material.

The combination of Roles and Instructions makes it possible to decompose a larger legal task into numbered Steps,

each an individual unit of Workflow execution. Changes in assignment, analytical function, or required output distinguish the Steps. A litigation analysis does not need to be performed through a single instruction asking the model to analyze an entire file. A workflow can first extract facts, then identify gaps, formulate legal issues, retrieve relevant legal rules, summarize those rules, apply them to the facts, identify risks, draft advice, and review the resulting reasoning. Each step has a defined function, and the output of one step can become an input for the next.

This decomposition is important because it makes legal work reusable. The expertise no longer exists only in the lawyer's ability to formulate a good prompt at a particular moment. It is captured in the Workflow: the Roles and Instructions applied in its Steps, the supporting Knowledge Cases, and the connections between tasks and their results. Roles and Instructions can also be shared across Workflows, making analytical approaches and assignments reusable beyond a single method.

Human involvement can also be designed into this sequence. A workflow may require a lawyer to provide information before a step begins, select between alternatives, exercise professional judgment, review an intermediate result, or validate the final output. Human involvement is therefore not external to the workflow. Where professional judgment or accountability requires it, it becomes part of the workflow design.

Knowledge and Client Cases

Roles and Instructions tell the Large Language Model what to do, but they do not by themselves provide the information necessary to do it. Legal work is information-intensive, and that information generally comes from two fundamen-

tally different sources: reusable legal knowledge and matter-specific information.

Knowledge Cases contain reusable knowledge that can support legal work across multiple matters. They may include legislation, case law, legal doctrine, precedents, regulatory guidance, model documents, internal know-how, checklists, previous analyses, and other forms of institutional expertise. Their defining characteristic is that the information is not limited to one individual client matter. It forms part of the knowledge base from which legal work can repeatedly draw.

Client Cases contain the information belonging to a particular client matter. This can include facts, contracts, correspondence, pleadings, evidence, meeting notes, emails, procedural documents, financial information, and other matter-specific material. The Client Case provides the factual and documentary environment within which the legal task must be performed.

Legal work commonly requires both. Consider a workflow designed to assess a contractual dispute. The relevant agreement, correspondence between the parties, chronology of events, and evidence of performance may come from the Client Case. The applicable legislation, relevant case law, internal precedents, and legal doctrine may come from Knowledge Cases. Neither information source is sufficient on its own. Legal knowledge without the client facts remains abstract, while client facts without the relevant legal framework cannot support a proper legal analysis.

This distinction also makes the information architecture reusable. A precedent stored as reusable knowledge does not need to be copied manually into every new matter. Conversely, confidential client information does not need to become part of a general knowledge repository merely because

it was used by a workflow. AI Workflows can draw on both information environments while preserving their different functions within the Legal AI System.

Retrieval-Augmented Generation

Having access to the necessary information creates another practical problem. A law firm may possess thousands or millions of pages of legislation, case law, correspondence, contracts, pleadings, internal know-how, and client documents. Providing all of this information to a Large Language Model for every task is neither necessary nor desirable.

Retrieval-Augmented Generation, commonly referred to as RAG, addresses this problem by retrieving information relevant to the task before the Large Language Model generates its output. Instead of expecting the model to work from everything that is available, the system identifies information that is relevant to the particular Workflow Step and provides that information as focused context.

The information needs of each Step are important. Different steps require different information. A fact extraction step may primarily require documents from the Client Case. A legal research step may primarily search Knowledge Cases. An application step may require both the factual output generated earlier in the workflow and the legal rules retrieved from relevant legal sources. A drafting step may additionally use an internal precedent to guide structure and terminology.

RAG therefore performs more than a technical search function. Within an AI Workflow, retrieval determines what information is brought into the working context for a specific task. The objective is not to maximize the quantity of information presented to the model, but to maximize its relevance.

This addresses context overload. An AI Workflow can use re-

trieval to select the information needed for a particular Step from the Client Case and its supporting Knowledge Cases, combining it with the relevant Role, Instruction, and preceding outputs. The model can then perform its task within a narrower and more purposeful context.

Retrieval does not eliminate the need for careful workflow design. The system must still determine what sources may be searched, what type of information is relevant, when retrieval should occur, and how retrieved information should be used. RAG is therefore not a substitute for legal process design. It is one of the mechanisms through which that design can be implemented; relevant information can also be supplied directly.

Conclusion

A Large Language Model provides a powerful general capability for processing and generating language, but a general capability is not yet a legal product. AI Workflows transform that capability into a reusable Legal Asset capturing an analytical method by addressing the three practical problems that arise when Large Language Models are applied to professional legal work.

Roles and Instructions address objective ambiguity by defining what the model is expected to do. Knowledge Cases and Client Cases address information deficiency by making the relevant reusable knowledge and matter-specific information available to the workflow. Retrieval-Augmented Generation can address context overload by selecting the information relevant to each particular workflow step. Human involvement can be incorporated wherever input, judgment, review, or validation is required.

This distinguishes AI Workflows from the Document Tem-

plates discussed in the previous chapter. Document Templates encode legal expertise through predefined Content and Logic applied with predominantly deterministic behaviour: given defined data and conditions, the system determines which content should appear and how the document should be assembled. AI Workflows operate probabilistically. They use Large Language Models to interpret and generate language where the appropriate output cannot be completely predetermined in advance.

The distinction does not make one mechanism superior to the other. They solve different forms of legal work and can also operate together. A Document Template may deterministically assemble the structure and standard provisions of a document while an AI Workflow drafts a fact-specific section. An AI Workflow may analyze a Client Case and produce structured information that subsequently populates a Document Template. Productizing legal expertise therefore requires the ability to determine which mechanism, or combination of mechanisms, is appropriate for the work being designed.

The challenge is consequently not simply to build templates or to write prompts. It is to identify suitable legal work, determine its requirements, and design the appropriate reusable legal assets in a systematic manner. Document Templates and AI Workflows require different implementation mechanisms, but both must emerge from a common design methodology.

The next chapter introduces that methodology and its three phases: Selection, Specification, and Design.

4 Design Methodology

The previous chapters introduced two complementary forms of reusable Legal Asset. Document Templates capture drafting knowledge in predefined Content and Logic, allowing established wording and drafting decisions to be applied consistently across matters. AI Workflows capture analytical methods in structured Steps that respond to each matter's information. Together, they allow legal expertise to be reused with less recurring lawyer effort, supporting more standardized and scalable legal work.

Creating such assets requires more than identifying a legal task and automating it. A useful Legal Asset must address recurring work, satisfy a clearly defined legal need, and be designed in sufficient detail for implementation. Without a structured methodology, teams risk developing assets without clear Objectives or moving directly from an initial idea to implementation without specifying what the asset must accomplish.

Legal AI Design provides the methodology for this transformation through three phases: Selection, Specification, and Design. Selection determines which work warrants investment. Specification defines what the selected assets must accomplish and how they will be evaluated. Design specifies how they should work. Together, these phases connect the business rationale for productization to a defined legal purpose and an implementation-ready design.

The V-model

The V-model connects this progression to validation. Its descending side moves from Opportunities through Requirements to Architecture and Design. After implementation, its ascending side evaluates the asset against the decisions that guided its development. Acceptance Criteria establish whether the Requirements are met, while Business Validation evaluates the business value realized through use of the implemented asset.

Selection connects Opportunity Identification on the definition side of the V-model with Business Validation on the validation side. Opportunity Identification establishes which Opportunities should proceed; Business Validation later evaluates the value realized through use of the implemented assets. Specification defines the Requirements and the Acceptance Criteria against which the resulting implementation will be assessed. Design develops Architecture and detailed specifications that make those Requirements implementable. The three phases therefore move from the reasons for developing an asset to the legal result expected and then to the method for producing it.

Legal AI Design V-Model

Phase	Definition	Validation
Selection	Opportunities (Chapter 5)	Business Validation (Chapter 6)
Specification	Requirements (Chapter 7)	Acceptance Criteria (Chapter 8)
Design	Architecture and Design (Chapters 9–12)	

The V-model gives each design decision a purpose and each

expected outcome a basis for evaluation. A design decision should respond to the Requirements, and the Requirements should express the Objective established for the selected Opportunity. Acceptance Criteria connect the implemented asset to its Requirements; Business Validation connects its use to the business rationale for capturing and reusing the expertise. These are distinct evaluations: satisfying the legal Requirements does not by itself establish how much business value an asset generates in practice.

Selection

Selection connects the identification of work with potential for reuse to validation of the business value realized after implementation. Opportunity Identification and Business Validation occupy the two ends of the V-model within this phase.

Opportunity Identification examines the firm's Legal Domains and Processes to reveal recurring work that could be captured in Document Templates and AI Workflows. Each Opportunity identifies one asset and states its Objective. A Process may involve both kinds of asset, with each recorded as its own Opportunity. Chapter 5 develops Opportunity Identification and its resulting inventory.

The inventory groups Opportunities by Legal Domain and Process, listing both kinds of asset together in descending order of frequency within each Process. This ordering expresses relative priority and provides a basis for selecting work that proceeds to Specification. Expected value is excluded from this prioritization because individual assets' value typically varies less significantly than their frequency of reuse.

Business Validation evaluates the business value realized af-

ter implementation. It assesses the revenue attributable to use of an asset or the cost reduction achieved compared with equivalent manual work, including assets supporting internal or non-billable work. Chapter 6 develops this evaluation and its measures.

The resulting Business Validation record connects the potential for reuse identified through the Opportunities inventory to the economic value generated in practice. Together, Opportunity Identification and Business Validation establish why recurring expertise warrants investment and whether its subsequent use delivers business value.

Specification

Once Opportunities have been selected, Specification defines what each asset must accomplish and how satisfactory performance will be assessed. Requirements Determination and Acceptance Criteria Setting provide the targets for Design and the basis for evaluating the resulting implementation.

Requirements translate the selected Objective into a clearly defined legal need. They describe the legal work to be performed, the information needed to perform it, and the characteristics expected of the result. This connects the rationale for developing the asset to the outcome it must produce, while leaving the mechanism for producing that outcome open. Chapter 7 develops Requirements Determination.

Acceptance Criteria establish the evidence needed to accept the result. They connect representative conditions of use to an expected result, giving the firm an agreed basis for deciding whether the implementation meets its Requirements. Defining them before Design keeps evaluation anchored in the intended legal outcome. Chapter 8 develops Acceptance Criteria Setting.

Requirements and Acceptance Criteria perform complementary functions. Requirements describe what the asset must accomplish; Acceptance Criteria establish how satisfactory performance will be recognized. Together, they make the selected Opportunity sufficiently explicit to guide Design without prescribing the design itself.

This distinction applies to both kinds of Legal Asset, while their evaluation reflects their different behaviour. Document Templates are predominantly deterministic, so their results can generally be assessed against predetermined drafting expectations. AI Workflows are probabilistic, so their criteria assess the substantive qualities an acceptable result must demonstrate while allowing variation in wording. The methodological structure remains the same in both cases.

Design

Design determines how the selected assets will produce the required results. It distinguishes Architecture from detailed Design. Architecture identifies shared components and maps their use across assets; detailed Design specifies the assets and components in sufficient detail for implementation.

Architecture reveals where several assets can rely on the same component and which assets will be affected when it changes. For Document Templates, it identifies shared Snippets and Questions. For AI Workflows, it identifies shared Roles and Instructions. Chapters 9 and 11 develop these respective forms of Architecture.

Document Template Design and Snippet Design translate drafting expertise into implementation-ready specifications. They keep the wording and structure connected to the rules governing how the document varies with the matter. Their deliverables combine the source Document with its associ-

ated components, preserving both the drafting content and the method needed to reproduce it. Chapter 10 develops this detailed design work.

AI Workflow Design organizes the analytical method into connected Steps, supported by Role Design and Instruction Design. Each Step applies an analytical approach and assignment, with supporting Knowledge Cases identified separately and matter information supplied during execution. The resulting specifications capture how the parts contribute to the intended result. Chapter 12 develops this detailed design work.

Architecture and detailed Design therefore perform complementary functions: one establishes shared structure and dependencies, while the other defines how the assets and their components operate. This supports consistent reuse across the portfolio and allows Document Templates and AI Workflows to contribute complementary drafting and analytical methods within a legal Process.

Conclusion

Legal AI Design connects the value sought from productization to the legal result expected and the method designed to achieve it. Selection identifies and prioritizes Opportunities and, after implementation, validates the business value realized through use. Specification defines Requirements and establishes Acceptance Criteria. Design develops shared component Architecture and implementation-ready specifications for the assets and their components.

After implementation, Acceptance Criteria establish whether the Requirements are met, while Business Validation evaluates the value realized through use. This completes the V-model's connection between development and evaluation,

keeping the finished asset tied to the reasons for developing it.

The following chapters develop each part of this methodology, beginning with Opportunity Identification: identifying recurring work that can be transformed into Document Templates and AI Workflows.

Part II

Selection

5 Opportunities

The foundations established in Part I explain how legal expertise can be captured in two complementary kinds of Legal Asset. Document Templates preserve drafting knowledge in predefined Content and Logic, while AI Workflows capture analytical methods in structured Steps that respond to each matter's information. Legal AI Design provides the methodology for transforming that expertise into reusable assets.

Part II develops Selection, which determines which recurring legal work warrants investment in productization. Opportunity Identification establishes where Document Templates and AI Workflows could capture recurring expertise and which Opportunities should proceed to Specification. Business Validation completes this connection after implementation by assessing whether the potential identified during Selection has translated into business value.

This chapter develops Opportunity Identification. It moves from the firm's Legal Domains and Processes to recurring documents and analytical tasks, defines an Objective for each Opportunity, and orders the resulting inventory by frequency within each Process. The inventory gives the firm a practical basis for selecting the work that proceeds to Specification.

Concept

An Opportunity identifies recurring legal work whose expertise could be captured in a reusable Document Template or

AI Workflow. It represents a prospect for productization, expressed through an Objective and situated within a Legal Domain and Process.

A Document Template Opportunity arises where established wording and drafting decisions can be captured for repeated use. It concerns a recurring document type whose Content and Logic can be defined in advance and applied to the circumstances of individual matters. An AI Workflow Opportunity arises where a recurring analytical task requires a method that responds to the information in each matter. It concerns legal work whose approach and assignments can be structured for repeated execution.

The two kinds of asset can contribute to the same Process. An AI Workflow may analyse information that informs drafting or review a document generated from a Document Template. Each remains a separate Opportunity because each captures a distinct unit of reusable expertise. Identifying them separately makes their intended purposes clear while allowing their relationship within the Process to remain visible.

Identifying an Opportunity does not commit the firm to developing it. Opportunity Identification first establishes where reusable expertise exists, then prioritizes the resulting Opportunities according to how frequently that expertise can be applied. A broad review prevents the firm from overlooking recurring work, while the ordering directs attention toward the work with the greatest potential for reuse.

Frequency provides the basis for prioritization because capturing expertise once has greater potential benefit when that expertise is used repeatedly. Expected value can also influence an Opportunity's attractiveness, but it is excluded from this prioritization: the value of individual Document Templates and AI Workflows typically varies less significantly

than the frequency with which they can be reused. Opportunity Identification therefore establishes relative priority through frequency rather than through a forecast of revenue, costs, or return on investment.

Each Opportunity also needs an Objective that describes what the asset is intended to accomplish. The Objective states the document to be generated or the analytical task to be performed in its legal context. It gives the Opportunity a clear purpose without specifying the Requirements, Architecture, or detailed Design through which that purpose will be achieved.

This distinction keeps the phases connected. Selection identifies and prioritizes the work; Specification develops the selected Objective into Requirements and Acceptance Criteria; Design determines how the resulting asset should work. After implementation, Business Validation assesses the business value realized through its use.

Framework

Opportunity Identification begins with the structure of the firm's practice and representative matters. Reviewing practice areas reveals the relevant Legal Domains; examining the work performed within them reveals recurring Processes, documents, and analytical tasks. The framework then makes each Opportunity explicit and orders it within its Process.

The framework consists of five steps.

Step 1 — Identify Legal Domains

A Legal Domain is a practice area or legal niche. It provides the highest-level grouping for the Opportunities inventory and preserves the substantive context in which expertise can be reused.

The appropriate level of granularity depends on the work being examined. A broad practice area may provide a useful Legal Domain where the work is relatively homogeneous. Where it contains substantially different types of legal work, narrower niches may make recurring expertise easier to identify. The purpose is to map the firm's actual legal work rather than reproduce its departmental structure.

Reviewing the firm's practice areas and representative matters helps establish this scope. Matters show how expertise is applied in practice and can reveal recurring work that a practice-area description alone would leave implicit. The result is a set of Legal Domains within which Processes and Opportunities can be identified consistently.

Step 2 — Identify Processes

A Process is a recurring sequence or category of legal work within a Legal Domain. It connects the substantive area of practice to the documents and analytical tasks that lawyers repeatedly produce or perform.

Depending on the practice, a Process may correspond to a formal legal procedure, a transaction, an advisory engagement, compliance work, or another recurring category of legal service delivery. Examining representative matters reveals the activities involved and the points at which lawyers repeatedly apply similar drafting decisions or analytical methods.

The Process provides a practical grouping for comparing Opportunities. A single Process can involve several recurring documents and analytical tasks, each with a different frequency of use. Keeping them within the same group makes it possible to identify which parts of that work offer the greatest potential for reuse.

At this step, the aim is to understand where recurring work

occurs. A Process should not be omitted simply because some of its work is complex or depends on professional judgment. The relevant question is which elements contain recurring expertise that could be captured in one of the two kinds of Legal Asset.

Step 3 — Identify Document Template and AI Workflow Opportunities

For each Process, identify the recurring document types and analytical tasks. These are the sources of Document Template and AI Workflow Opportunities.

The analysis of documents should focus on recurring types rather than individual files. A document produced for one matter provides evidence of drafting knowledge; the Opportunity lies in capturing the wording and decisions that can be applied across comparable matters. Where known circumstances can determine the document's wording and structure, a Document Template can preserve those choices in reusable Content and Logic.

Analytical tasks provide the corresponding source of AI Workflow Opportunities. Lawyers repeatedly analyse information, assess legal issues, review materials, extract relevant facts, compare information, classify matters, and prepare legal reasoning. Where a recurring task requires contextual analysis, an AI Workflow can capture the method while allowing the result to respond to the particular matter.

The distinction rests on the expertise being captured. A Document Template applies drafting choices settled in advance; an AI Workflow applies an analytical method to matter information. A task that produces written output may therefore be an AI Workflow Opportunity where its substance depends on contextual analysis. Where the work also involves generating a document from established drafting choices, the Process

may contain Opportunities for both kinds of asset.

The result is a set of candidate Document Templates and AI Workflows, each connected to the Legal Domain and Process in which its expertise can be reused.

Step 4 — Specify Objectives

Specify an Objective for each Opportunity so that its intended purpose is clear. The Objective identifies the legal work to be productized and gives the later Requirements a defined starting point.

For a Document Template, the Objective follows the structure:

Generate [document type] for [legal context]

This formulation identifies the document to be generated and the legal circumstances in which it will be used. It states the purpose of the proposed Document Template without prescribing its Content, Logic, or component structure.

For an AI Workflow, the Objective follows the structure:

[verb] [noun] for [legal context]

The verb and noun identify the analytical task, while the legal context establishes where that task applies. The wording should make the intended work understandable without prescribing the Steps or other components through which it will be performed.

For both kinds of asset, keep the Objective concise and focused on purpose. Detailed descriptions of inputs, substantive output expectations, and evaluation belong in Specification. Choices about components, prompts, models, integrations, and execution belong in Design or implementation. Making the Objective clear at this stage allows those later

decisions to remain anchored in the legal work selected for productization.

Step 5 — Prioritize by Frequency

Within each Process, list Document Template and AI Workflow Opportunities together in descending order of frequency. The more often the underlying document is needed or analytical task is performed, the greater the potential to reuse the captured expertise and the higher the Opportunity's relative priority.

The review of representative matters provides a practical basis for comparing recurrence. Consider how often each document or task arises across the work within scope, rather than assuming that every Opportunity occurs once in every matter. Some documents or tasks may be required throughout a Process, while others arise only under particular circumstances.

Both kinds of asset belong in the same ordered list because frequency of reuse is the common basis for prioritization. Separate lists for Document Templates and AI Workflows would obscure their relative priority within the Process. The ordering applies within each Process; it does not establish a single ranking across all Legal Domains and Processes.

Position in the inventory expresses relative priority. Expected value is excluded from this ordering because individual assets' value typically varies less significantly than their frequency of reuse. The method therefore does not require a financial assessment before an Opportunity can be placed in the inventory.

The resulting order gives the firm a practical basis for selecting which Opportunities should proceed to Specification. Selection establishes the work to take forward; Business Val-

idation later assesses the value realized through use of the implemented assets. Anticipated recurrence at this stage and actual use after implementation serve different purposes within the V-model.

Resulting Deliverable

The resulting deliverable is the Opportunities inventory: a structured list grouped by Legal Domain and Process. Within each Process, Document Template and AI Workflow Opportunities appear together in descending order of frequency. Each entry identifies the kind of Legal Asset and states its Objective.

The grouping preserves the legal and operational context without repeating it in a separate profile for every Opportunity. The entry makes the proposed asset's purpose explicit, while its position communicates relative priority. The inventory takes the following form:

Opportunities

[legal domain]

[process]

[Document Template or AI Workflow] — [objective]

[...]

[...]

[...]

This deliverable brings identification and prioritization together. It shows where reusable expertise exists, what each proposed asset should accomplish, and which Opportunities

have greater potential for repeated use within their Process. Selected Opportunities proceed to Requirements Determination and Acceptance Criteria Setting, where their intended legal results and the basis for evaluating them are defined.

The inventory is distinct from the Business Validation record. The former establishes potential and relative priority before Specification; the latter records Value per Use, Annual Frequency, and Annual Value after implementation. Keeping these purposes clear connects the reasons for developing an asset to the business results subsequently achieved through its use.

Conclusion

Opportunity Identification establishes the starting point for Selection by reviewing the firm's Legal Domains and Processes, identifying recurring documents and analytical tasks, and defining an Objective for each proposed Document Template or AI Workflow. Listing both kinds of asset together in descending order of frequency within each Process makes their relative priority explicit.

The Opportunities inventory provides a practical basis for selecting the work that proceeds to Specification. Requirements and Acceptance Criteria then define what the selected assets must accomplish and how satisfactory performance will be assessed, before Architecture and detailed Design specify how they should work.

Chapter 6 develops Business Validation, the corresponding evaluation on the ascending side of the V-model. After implementation, it assesses whether use of the assets has translated the potential identified during Selection into realized business value.

6 Business Validation

Chapter 5 identified recurring legal work whose expertise could be captured in Document Templates and AI Workflows. It defined an Objective for each Opportunity and ordered the inventory by frequency within each Process, providing a practical basis for selecting the work that proceeds to Specification. Business Validation completes this connection after implementation by assessing whether the potential identified during Selection has translated into business value.

An asset's suitability for reuse does not by itself establish the economic contribution it makes in practice. A Document Template may generate documents quickly but be used only occasionally. An AI Workflow may support frequent internal analysis without being sold as a separate service. Business Validation therefore considers both the value realized through each use and how often the asset is actually used.

This chapter develops the validation side of Selection in the V-model. Although presented alongside Opportunity Identification, Business Validation takes place after the selected assets have been specified, designed, implemented, and used. Its result is a record of the business value generated through repeated use of the captured expertise.

Concept

Business Validation establishes whether reusable legal expertise delivers business value in practice. Legal Assets can

contribute to revenue-generating services or reduce the effort required for internal or non-billable work. Their economic value is therefore based on the revenue attributable to their use or the cost reduction achieved compared with equivalent manual work.

The same approach applies to Document Templates and AI Workflows. A Document Template captures established drafting decisions so that they can be applied repeatedly across matters. An AI Workflow captures an analytical method that responds to each matter's information. Both can reduce recurring lawyer effort, but the business value realized depends on how the resulting asset is used.

For revenue-generating work, the relevant value is the revenue attributable to use of the asset. A Document Template may support a fixed-fee drafting service, while an AI Workflow may contribute to a paid review or analysis. Where the asset forms part of a broader service, the assessment must identify the revenue attributable to its contribution rather than assign the entire service revenue to every asset involved.

Reduced production effort does not inherently eliminate the commercial value of legal expertise. A document or analysis can remain valuable to a client even when the asset allows it to be produced more efficiently. Pricing and the firm's value proposition can therefore affect the revenue realized through use. Business Validation records that realized contribution; it does not prescribe a price derived from the cost of automated production or require every asset to be sold separately.

For internal or non-billable work, the relevant value can instead be the cost reduction achieved compared with equivalent manual work. An asset that reduces the effort required for internal drafting, knowledge work, or analysis can create business value even where no client is charged for its use.

The comparison should concern work of equivalent scope and quality, including any lawyer review or completion effort still required when the asset is used.

Value per Use records this economic value each time an asset is used, based on attributable revenue or resulting cost reduction. Annual Frequency records how often the asset is used during the year. Together, they establish Annual Value:

$$\text{Annual Value} = \text{Value per Use} \times \text{Annual Frequency}$$

These measures preserve the connection between value and reuse. An asset with a modest Value per Use can generate substantial Annual Value through frequent application. An asset with a high Value per Use may make a smaller annual contribution if it is rarely used. Annual Value provides a common basis for evaluating the business value of implemented Document Templates and AI Workflows, including those supporting internal work.

Business Validation has a different purpose from Opportunity Identification and Acceptance Criteria. Opportunity Identification prioritizes potential reuse through frequency within each Process; expected value is excluded from that prioritization. Acceptance Criteria establish whether the implemented asset meets its Requirements. Business Validation evaluates the economic value realized through use. Meeting legal Requirements and generating business value are related outcomes, but each requires its own evaluation.

Framework

Business Validation brings together three measures for each implemented asset: Value per Use, Annual Frequency, and Annual Value. The assessment begins with the Document Template or AI Workflow and its Objective, preserving the connection to the Opportunity identified during Selection.

Step 1 — Determine Value per Use

The first step is to establish the economic value realized each time the asset is used. The basis is the revenue attributable to that use or the cost reduction achieved compared with equivalent manual work, according to the asset's contribution.

For revenue-generating services, evidence can come from the fees actually charged for work supported by the asset. If a Document Template generates the document delivered through a separately priced service, the attributable revenue can be established from that service. If an AI Workflow supports one part of a larger engagement, the assessment should use a defensible allocation of revenue to that contribution. The same principle applies to either kind of Legal Asset.

The price previously charged for equivalent manual work can provide commercial context, but it does not establish the revenue realized by the implemented asset. Productization may change the service's price, demand, or delivery model. Business Validation should therefore use evidence of the revenue actually attributable to use rather than a proposed price or an assumption that every use results in a sale.

For work assessed through cost reduction, compare the cost of performing the equivalent task manually with the cost of performing it using the asset. The comparison should include the effort and operating costs relevant to delivering the same result. For example, if manual work costs 120 monetary units and equivalent work using the asset costs 40, the resulting Value per Use is 80 monetary units.

A reduction in drafting or analysis time must be considered alongside the work still needed to complete the task. Required review, correction, and completion effort affects the cost reduction actually achieved. This keeps the measure connected to the delivered legal result rather than to the speed

of generating an initial output alone.

The assessment should make clear which basis supports the recorded amount. Attributable revenue and cost reduction describe different economic contributions; the same benefit should not be counted twice. Where the value varies across uses, an average supported by the observed work can provide the Value per Use for the period assessed.

Step 2 — Record Annual Frequency

The second step is to record how often the Document Template or AI Workflow is used during the year. Annual Frequency concerns actual use of the implemented asset, including uses supporting internal or non-billable work.

Historical manual volume helped reveal the potential for reuse during Opportunity Identification. Actual use after implementation shows how much of that potential has been realized. Productization can affect demand and working practices: faster delivery, more predictable pricing, or easier access may increase use, while limited adoption may leave the asset used less often than the original manual process.

Annual Frequency should therefore be supported by records of use rather than copied from the Opportunity's anticipated frequency or treated as a forecast of future sales. A consistent definition of a use should connect the count to the contribution measured in Value per Use. For a Document Template, this may be the generation of a document for a matter; for an AI Workflow, it may be the execution of the defined analytical task. Technical retries or intermediate Steps should not be counted as separate completed contributions where Value per Use refers to the resulting document or overall task.

The value measure and the use count should cover the same period and scope. If an asset has been in use for less than a

full year, the available evidence should be identified as partial-year evidence. An annual projection can inform planning, but it should remain distinguishable from the business value already realized.

Step 3 — Calculate Annual Value

The third step is to multiply Value per Use by Annual Frequency:

$$\text{Annual Value} = \text{Value per Use} \times \text{Annual Frequency}$$

The formula applies to both kinds of Legal Asset and to either basis of value. For an asset with attributable revenue of 200 monetary units per use and 150 uses during the year, Annual Value is 30,000 monetary units. For an internal asset achieving a cost reduction of 80 monetary units per use over 150 uses, Annual Value is 12,000 monetary units.

The calculation provides a simple, common measure of the value generated through repeated use. Its interpretation follows the basis used for Value per Use: a revenue-based amount records attributable revenue, while a cost-reduction amount records the reduction achieved against equivalent manual work. Annual Value is not, by itself, a calculation of net profit or a return on development investment.

Reviewing the result connects the implemented asset to the business rationale established during Selection. Limited Annual Value may reflect low use, a small economic contribution per use, or both. Distinguishing these factors helps the firm understand whether the captured expertise is delivering the anticipated benefits and where changes to the asset or its use may be useful. This evaluation informs subsequent decisions without replacing the frequency-based prioritization of Opportunities.

Resulting Deliverable

The resulting deliverable is a Business Validation record for each implemented Document Template or AI Workflow. The record identifies the asset by its type and Objective and brings together Value per Use, Annual Frequency, and Annual Value. Both kinds of asset use the same structure, providing a common basis for evaluating the business value generated through repeated use of the captured expertise.

A generic representation of the record is:

Business Validation

[Document Template or AI Workflow] — [objective]

Value per Use

[amount]

Annual Frequency

[uses per year]

Annual Value

[amount]

Supporting evidence should make the basis of Value per Use and the period and scope of the use count understandable. It accompanies the three measures so that the result can be interpreted consistently, while the record retains the same structure for revenue-generating and internal uses.

The Business Validation record evaluates an implemented asset. The Opportunities inventory records potential reuse and relative priority before Specification. Keeping these deliverables distinct allows the firm to connect the reasons

for developing an asset to the business results subsequently achieved through its use.

Conclusion

Business Validation completes the connection between the potential identified during Selection and the business value realized after implementation. Value per Use captures attributable revenue or cost reduction, Annual Frequency captures actual use during the year, and Annual Value brings the two together. The resulting record allows both Document Templates and AI Workflows to be evaluated on a common basis.

This completes the treatment of Selection in Part II, covering Opportunity Identification on the descending side of the V-model and Business Validation on its ascending side. In the development sequence, selected Opportunities proceed from the ordered inventory to Specification; their business value is validated once the resulting assets have been implemented and used.

Chapter 7 develops Requirements Determination. It translates each selected Objective into the Legal Task, Input Scope, and Output Characteristics. Acceptance Criteria then establish how satisfactory performance will be assessed, before Architecture and detailed Design specify how the assets should work.

Part III

Specification

7 Requirements

Part III develops Specification, which defines what each selected Legal Asset must accomplish and how satisfactory performance will be assessed. Opportunity Identification established where Document Templates and AI Workflows could capture recurring expertise and ordered those Opportunities by frequency within each Process. Selected Opportunities now proceed to Specification. Business Validation evaluates the value realized through use of the resulting assets after implementation.

Specification comprises Requirements Determination and Acceptance Criteria Setting. This chapter develops Requirements Determination: it translates each selected Objective into the Legal Task, the circumstances in which it will be performed, the Input Scope needed to complete it, and the Output Characteristics that a satisfactory result must demonstrate. Chapter 8 establishes the Acceptance Criteria against which the implementation will be evaluated.

The resulting Requirements record connects the value sought through productization to a defined legal outcome. It gives Architecture and detailed Design a clear target while leaving open the mechanism through which that outcome will be produced.

Concept

An Objective and Requirements serve different functions within Legal AI Design. The Objective identifies the legal work to be productized at a level appropriate for selecting an Opportunity. Requirements develop that purpose into a sufficiently precise description of what the resulting Document Template or AI Workflow must accomplish.

The Objective remains the starting point. For a Document Template, it takes the form “Generate [document type] for [legal context]”; for an AI Workflow, it takes the form “[verb] [noun] for [legal context]”. Requirements make the intended work explicit by defining the Legal Task, Input Scope, and Output Characteristics under that Objective.

These three elements answer related questions. The Legal Task defines the work and the circumstances in which it is required. The Input Scope identifies the matter information, legal knowledge, and other inputs necessary to perform it. The Output Characteristics establish what a satisfactory result must contain or demonstrate. Together, they describe the intended legal outcome at the level needed to guide Design and establish Acceptance Criteria.

Requirements define what the asset must accomplish while leaving decisions about how it should work to Design. They may establish that particular information is needed, that specified content must appear, or that a conclusion must be supported by the available sources. The component structure, automation rules, analytical Steps, and implementation mechanisms through which those outcomes are achieved are developed later.

For a Document Template, Requirements describe the document generation needed and the essential characteristics of the resulting document. Established drafting decisions will

subsequently be captured in Content and Logic, with Document Automation applying them through predominantly deterministic behaviour. Requirements provide the target for those decisions without yet specifying the Document, Fields, Questions, or Snippets through which they will be implemented.

For an AI Workflow, Requirements describe the legal work to be performed or supported, the information it requires, and the qualities of its result. Design will subsequently organize the analytical method into Steps, each applying a Role and an Instruction, with supporting Knowledge Cases identified where needed. Requirements establish the purpose and boundaries of that method before its components are designed.

AI Workflows also require attention to the probabilistic behaviour of Large Language Models. Equivalent inputs need not produce identical wording. Precise Requirements remain possible because they define the substantive characteristics that an acceptable result must demonstrate. Different formulations can satisfy the same Requirements where they address the required issues, reflect the relevant information, and support the intended legal outcome.

Requirements also have a distinct place in the V-model. They establish what the asset must accomplish on its descending side; Acceptance Criteria provide the basis for evaluating whether the implementation meets those Requirements on its ascending side. Business Validation addresses the separate question of whether use of the asset realizes business value. Defining the required legal result therefore supports validation without treating legal performance and business value as the same assessment.

Framework

Requirements Determination proceeds through three steps for each selected Document Template or AI Workflow. The selected Objective provides the common reference throughout: each step develops an aspect of the capability needed to fulfil it.

Step 1 — Define the Legal Task

Define the legal work that the asset must perform or support and the circumstances in which that work is required. The Legal Task makes the Objective sufficiently precise to establish the scope of the asset.

For a Document Template, the Legal Task describes the document generation to be supported. It identifies the kind of document to be produced and the legal circumstances in which it will be used. Where different circumstances materially affect the required document, the description should make those boundaries clear without prescribing the automation rules that will implement them.

For an AI Workflow, the Legal Task defines the analytical work to be performed or supported. This may involve analysing, reviewing, extracting, comparing, classifying, or preparing legal information. The description should explain what the work must achieve in its legal context, leaving the sequence of Steps and the allocation of Roles and Instructions to Design.

The Legal Task should be specific enough to prevent different designers from assuming different intended capabilities. A broad Objective may leave questions about which matters, issues, or activities are included. Resolving those questions here establishes a common understanding of the work and limits implicit scope expansion during Architecture and detailed Design.

The result is a bounded description of the required legal activity and its circumstances. It provides the basis for identifying the information needed to perform that activity and the characteristics its result must demonstrate.

Step 2 — Define the Input Scope

Define the information needed to complete the Legal Task. The Input Scope identifies the relevant matter information, legal knowledge, and other necessary inputs, making explicit the informational basis on which the asset is expected to operate.

For a Document Template, the Input Scope identifies the relevant Client Case information and any additional information needed to generate the document. This may include the parties' details, relevant dates, agreed terms, factual circumstances, and choices that affect the required document. The description establishes the categories of information that must be available without yet defining individual Fields or Questions.

Design subsequently determines how that information is used. Fields identify variable information to be inserted into the Document, Questions request required user input, and Logic governs inclusion, selection, calculation, and repetition. Requirements establish the information need that those design decisions must address.

For an AI Workflow, the Input Scope identifies the matter information and supporting legal knowledge needed for the analysis, including relevant information from Client Cases and Knowledge Cases where applicable. It may also identify other necessary inputs, such as material produced by related legal work. The description should make clear which information the method needs and the scope of the knowledge on which it must rely.

Identifying the required legal knowledge does not yet allocate particular Knowledge Cases to Workflow Steps. That allocation belongs to AI Workflow Design, which connects supporting knowledge to each Step's Role and Instruction. The Input Scope first establishes what information is necessary for the overall Legal Task.

The scope should be complete enough to support the required work while remaining independent of the mechanisms used to obtain or process the information. Retrieval, data entry, information transfer, and other implementation decisions follow from the need defined here. Information outside the stated scope should not silently become a dependency; if additional inputs prove necessary, the Requirements should be revised accordingly.

Step 3 — Define the Output Characteristics

Define what a satisfactory result must contain or demonstrate. The Legal Task establishes the work to be performed, and the Input Scope establishes its informational basis. The Output Characteristics give Design a clear target for the resulting legal work product.

For a Document Template, these characteristics describe the generated document. They may address required content, the treatment of relevant circumstances, document structure, and necessary presentation qualities. Where particular circumstances require different content, the Requirements should state the resulting document's required characteristics. The Content, Logic, and formatting that produce those results are specified during detailed Design.

This allows drafting expectations to be precise without prematurely designing the template. A requirement that the document reflect the applicable terms, for example, establishes an outcome. Selecting the wording, defining the conditions

under which it appears, and determining whether a shared Snippet should contain it are later design decisions.

For an AI Workflow, the Output Characteristics define the substantive qualities of the required analysis or other legal result. Depending on the Legal Task, they may identify issues to be addressed, information to be extracted, distinctions to be made, sources to be connected to findings, or uncertainties to be acknowledged. They should describe what satisfactory completion of the work requires rather than rely on broad statements that the output must be “correct” or “professional”.

These characteristics should allow variation in wording where it is compatible with the required legal outcome. Different executions may express or organize the analysis differently while satisfying the same Requirements. Variation remains acceptable only where the result demonstrates the substantive qualities defined for the Legal Task.

Output Characteristics provide the basis for Acceptance Criteria. The next stage of Specification pairs Test Conditions with an Expected Result for each criterion, establishing the representative inputs and legal circumstances under which the asset will be assessed and what the evaluator must observe. Requirements first define the intended result; Acceptance Criteria establish the evidence needed to accept it.

Together, the three steps produce a coherent statement of the required capability. The Legal Task defines the work, the Input Scope identifies what is needed to perform it, and the Output Characteristics define satisfactory completion. Reviewing them together helps reveal gaps, such as an expected result that depends on information absent from the Input Scope or a characteristic that falls outside the Legal Task.

Resulting Deliverable

The deliverable is a Requirements record for each selected Document Template or AI Workflow. It retains the asset type and Objective from the Opportunities inventory as its heading and brings together three elements: Legal Task, Input Scope, and Output Characteristics.

The record uses the same structure for both kinds of Legal Asset. Its substance reflects the document generation or analytical work concerned, while its common form makes the intended outcome clear before Architecture and detailed Design begin.

Requirements

[Document Template or AI Workflow] — [objective]

Legal Task

[description]

Input Scope

[description]

Output Characteristics

[description]

For a Document Template, the Legal Task describes the required document generation and its circumstances, the Input Scope identifies the necessary matter information and other inputs, and the Output Characteristics describe the document that must result. For an AI Workflow, the same elements describe the analytical work, its informational basis, and the substantive qualities the result must demonstrate.

The heading preserves the connection to the selected Opportunity. The three elements develop that Opportunity into a defined legal purpose and make its intended outcome explicit. The record can therefore guide the development of the asset while remaining focused on what it must accomplish.

Conclusion

Requirements Determination translates each selected Objective into a Legal Task, Input Scope, and Output Characteristics. The resulting Requirements record defines the intended legal outcome and gives Design a clear target, while preserving the distinction between the required result and the method used to produce it.

Chapter 8 develops Acceptance Criteria Setting. Each criterion pairs Test Conditions with an Expected Result to establish how satisfactory performance will be assessed. Together, Requirements and Acceptance Criteria complete Specification, providing the targets and evaluation standards that guide Architecture and detailed Design and support validation after implementation.

8 Acceptance Criteria

Specification defines what each selected Legal Asset must accomplish and how satisfactory performance will be assessed. Chapter 7 developed Requirements Determination, translating each selected Objective into the Legal Task, Input Scope, and Output Characteristics. This chapter develops Acceptance Criteria Setting: establishing the evidence needed to accept the resulting implementation.

Requirements describe the intended result; Acceptance Criteria establish how that result will be evaluated. Each criterion pairs Test Conditions with an Expected Result, specifying the representative inputs and legal circumstances under which the asset is assessed and what the evaluator must observe. Together, Requirements and Acceptance Criteria provide the targets and evaluation standards that guide Architecture and detailed Design.

The criteria are defined during Specification and applied after implementation. In the V-model, they provide the basis for evaluating whether the implemented Document Template or AI Workflow meets its Requirements. Business Validation addresses the separate question of whether use of the asset realizes business value.

Concept

Acceptance Criteria allow quality to be assessed against agreed expectations. A Requirements record defines the

work through its Legal Task, identifies the necessary information through its Input Scope, and describes a satisfactory result through its Output Characteristics. Acceptance Criteria make those expectations testable under defined circumstances.

This distinction matters because a description can be clear enough to guide Design while leaving uncertainty about what evidence would establish satisfactory performance. An Output Characteristic might require a generated document to include the applicable payment deadline or an analysis to identify the relevant contractual obligations. Acceptance Criteria specify the inputs and circumstances used to assess those expectations and the result that must be observed.

Each criterion therefore connects a testing situation to an observable result. The Test Conditions establish the representative inputs and legal circumstances. The Expected Result states what the evaluator must find in the resulting document or analysis. Reviewing the criteria alongside the Requirements record makes clear which expectations they assess and reveals any material expectations that remain untested.

The criteria reflect how each kind of Legal Asset produces its results. Document Templates capture established drafting decisions in Content and Logic. Document Automation applies those decisions through predominantly deterministic behaviour, so criteria can address the generated document's content, drafting logic, and formatting. Structure, calculations, numbering, and references may also be assessed where they form part of the required result.

A required clause, for example, can be assessed under circumstances in which it must appear and circumstances in which it must be omitted. A required calculation can be assessed by supplying defined values and comparing the

amount produced with the expected amount. Required document structure can be assessed by checking that specified sections appear in the appropriate sequence. These observations evaluate the outcome of the drafting method without prescribing its implementation.

AI Workflows capture analytical methods applied by Large Language Models with probabilistic behaviour. Their conclusions and wording can vary in response to the matter's information, and equivalent inputs need not produce identical formulations. Acceptance Criteria therefore define the substantive qualities an acceptable result must demonstrate while allowing variation in wording where it is compatible with the Requirements.

For an AI Workflow analysing contractual obligations, those qualities might include identifying the obligations represented in the test material, distinguishing the parties to whom they apply, connecting findings to the relevant sources, and acknowledging uncertainty where the information does not support a conclusion. Different formulations can satisfy the same criterion when they demonstrate the required qualities. Permissible variation does not remove the need for a precise basis for acceptance.

Both kinds of asset can thus be assessed against their Requirements using standards suited to their behaviour. Defining those standards before Design gives the firm a practical basis for deciding whether the completed implementation fulfils its legal purpose.

Framework

Acceptance Criteria Setting proceeds through three steps for each selected Document Template or AI Workflow. First identify the Acceptance Criteria from its Requirements record.

Then, for each criterion, define the Test Conditions and the Expected Result. The asset's Objective remains the common reference connecting the criteria to the selected Opportunity and its Requirements.

Step 1 — Identify the Acceptance Criteria

Review the Legal Task, Input Scope, and Output Characteristics to identify the aspects of satisfactory performance that need to be assessed. Each Acceptance Criterion should address a specific expectation grounded in the Requirements. Give it a descriptive name that makes the focus of the assessment clear.

For a Document Template, criteria may address required content, the treatment of particular legal circumstances, calculations, or formatting. For an AI Workflow, they may address the identification of relevant issues, attribution of obligations, support for findings, or acknowledgement of uncertainty. Select the criteria according to the required legal outcome.

One expectation may require several criteria where materially different circumstances call for different results. Conversely, related observations may belong together where they establish whether one expectation has been met. There need not be one criterion for each statement in the Requirements record; the criteria together should cover all material expectations at a level that permits clear assessment.

The result of this step is a set of named criteria, each with a defined focus. The following steps make each criterion testable by pairing representative Test Conditions with an observable Expected Result. If developing that pair reveals distinct assessment needs, refine or separate the criteria accordingly.

Step 2 — Define the Test Conditions

Define the representative inputs and legal circumstances under which the asset will be assessed. The Legal Task establishes the work and its boundaries, the Input Scope identifies the information available, and the Output Characteristics indicate which aspects of the result need to be observed. Test Conditions turn these descriptions into a sufficiently specific situation for evaluation.

Depending on the expectations being assessed, the conditions may include particular input values, available source documents, factual circumstances, agreed terms, user choices, or other relevant information. They should make clear what is supplied and which circumstances apply, so that an evaluator can interpret the result against a common testing situation.

For Document Templates, Test Conditions often identify the matter information and choices that affect document generation. These may include a party type, a monetary amount, a relevant date, or circumstances requiring an optional provision. Where the expected document differs according to those inputs, the conditions should specify the combination being assessed.

Materially different circumstances may require separate criteria. If a provision must appear when a condition is present and be omitted when it is absent, testing only its inclusion would leave the omission unassessed. Separate criteria can establish the evidence needed for both outcomes. Calculations can similarly require more than one scenario where different ranges of input affect the required result.

For AI Workflows, Test Conditions identify the matter information, source material, and legal or factual circumstances relevant to the analysis. These may include information from

Client Cases and supporting legal knowledge from Knowledge Cases where applicable. The conditions establish the informational basis for assessment without allocating Roles, Instructions, or Knowledge Cases to particular Steps; those decisions belong to Design.

Different information may call for different substantive results. One criterion might assess an analysis where the supplied material establishes an obligation, while another assesses a situation where the material is incomplete or ambiguous. If the Requirements call for uncertainty to be acknowledged, the latter situation provides evidence of whether the implemented Workflow does so.

The aim is to cover material distinctions in the Requirements through representative conditions. The criteria need not enumerate every conceivable factual variation. Their coverage should be sufficient to assess the intended Legal Task, using information within the Input Scope, against the required Output Characteristics.

Step 3 — Define the Expected Result

Define what the evaluator must observe under the Test Conditions. The Expected Result establishes the evidence needed to accept the result produced in that situation. It should be precise enough to support a decision about whether the criterion is met.

For Document Templates, the Expected Result can specify required content, the presence or absence of a clause, the sequence of sections, a calculated value, or formatting and numbering. Where the criterion assesses drafting logic, it states the observable document outcome under the defined circumstances. The rules that produce that outcome are subsequently specified during Design.

A statement that a document must be “correct” or “professional” provides little basis for acceptance because reviewers may interpret it differently. A statement that a specified section must follow another section under the defined conditions provides an observable standard. Similarly, an expected calculation should state the result required for the supplied values, rather than merely require an accurate calculation.

For AI Workflows, the Expected Result specifies the substantive qualities the output must demonstrate. Depending on the Requirements, it may identify issues to be addressed, information to be extracted, distinctions to be made, sources to be connected to findings, or uncertainties to be acknowledged. The defined test material should give the evaluator a concrete basis for assessing those qualities.

A criterion concerning contractual obligations, for example, can specify which obligations in the supplied material must be identified and attributed to each party. Where source support is required, it can establish which provisions must support the findings. If the material leaves an issue unresolved, the Expected Result can require that limitation to be expressed without presenting an unsupported conclusion as established.

Exact wording is appropriate only where the Requirements call for it. Otherwise, the Expected Result should allow different formulations that demonstrate the same substantive qualities. Precision comes from specifying what the result must establish, not from requiring every acceptable analysis to reproduce a predetermined paragraph.

One Expected Result may describe several related observations needed to assess a criterion. Where materially different conditions or independently assessable outcomes would make the criterion unclear, separate criteria make the basis

for acceptance easier to apply. Their number follows the expectations being assessed rather than a fixed rule.

Review the completed criteria against the Legal Task, Input Scope, and Output Characteristics. Each criterion should have a clear basis in the Requirements, and the criteria together should provide evidence for all material expectations. If a proposed criterion introduces a new capability or requires information outside the Input Scope, resolve that issue in the Requirements before using it as a basis for acceptance.

Resulting Deliverable

The deliverable is an Acceptance Criteria record for each selected Document Template or AI Workflow. Its heading retains the asset type and Objective used in the Opportunities inventory and Requirements record. Under that heading, each named criterion pairs Test Conditions with an Expected Result.

The same structure applies to both kinds of Legal Asset. For Document Templates, its substance can address content, drafting logic, and formatting. For AI Workflows, it defines the substantive qualities an acceptable result must demonstrate while allowing appropriate variation in wording.

[Document Template or AI Workflow] — [objective]

[acceptance criterion]

Test Conditions

[description]

Expected Result

[description]

[...]

The heading preserves the connection to the selected Opportunity and its defined legal purpose. The criteria make the evidence for acceptance explicit, so the record can guide Design and later evaluation. Read alongside the Requirements, it establishes what satisfactory performance must demonstrate without prescribing the component structure or implementation mechanism.

Conclusion

Acceptance Criteria Setting completes Specification by establishing how satisfactory performance will be assessed. Each criterion pairs Test Conditions with an Expected Result, connecting representative inputs and legal circumstances to the evidence needed for acceptance. The resulting record provides standards suited to the predominantly deterministic behaviour of Document Templates and the probabilistic behaviour of AI Workflows.

Requirements and Acceptance Criteria together guide Architecture and detailed Design and provide the basis for evaluating the implementation against its intended legal outcome. Business Validation subsequently evaluates the busi-

ness value realized through use of the implemented asset, maintaining the connection to Selection on the validation side of the V-model.

Part IV develops Document Automation Design. Chapter 9 begins with Document Template Architecture, comparing the selected templates' Requirements to identify shared Snippets and Questions and map their dependencies. Detailed Document Template Design and Snippet Design then specify how the required results will be produced. Legal Intelligence Design follows the same progression from Requirements to Architecture and detailed Design for AI Workflows.

Part IV

Document Automation Design

9 Document Template Architecture

Selection identified which Opportunities should proceed to Specification. Requirements Determination then defined each selected Legal Asset's Legal Task, Input Scope, and Output Characteristics, while Acceptance Criteria Setting established the Test Conditions and Expected Results against which its implementation will be assessed. Together, these decisions establish what each asset must accomplish and how satisfactory performance will be evaluated.

Part IV develops Document Automation Design. For Document Templates, Design begins with Document Template Architecture: identifying shared Snippets and Questions and mapping them to the templates and Snippets that use them. Detailed Document Template Design and Snippet Design then specify how those components will operate. Architecture provides the common structure through which drafting decisions and information needs can be addressed consistently across documents.

This connection matters for productization. A Document Template captures established drafting decisions in reusable Content and Logic. Where several templates require the same drafting or user input, shared components allow that expertise to be developed and maintained together. Document Template Architecture makes those opportunities for reuse and their dependencies explicit before detailed design be-

gins.

Concept

Requirements describe the intended legal result while leaving the mechanism for producing it open. The Legal Task establishes the work to be performed, the Input Scope identifies the necessary information, and the Output Characteristics describe what a satisfactory document must contain or demonstrate. Acceptance Criteria provide the corresponding basis for evaluation after implementation.

Document Template Architecture compares the selected templates' Requirements to identify recurring drafting decisions and user input needs. Recurring Content, Logic, and formatting can become reusable Snippets. Questions identify the user input needed by the templates and their Snippets. Mapping these components to their uses shows which documents depend on the same decisions and which components must be designed together.

The Architecture therefore connects Specification to detailed Design. Its focus is the shared components and their dependencies. The complete wording, drafting rules, formatting, variable Fields, and use of answers are specified subsequently in Document Template Design and Snippet Design.

Snippets are reusable components containing Content, Logic, and formatting. They capture a coherent part of the firm's drafting method that can be used wherever the same legal and documentary needs recur. A Snippet may represent a clause, definition, party identification block, signature block, notice, or standard explanatory passage.

Content supplies the predetermined text. Logic specifies the rules governing its inclusion, selection, calculation, and repetition. Formatting preserves the presentation needed for

the resulting document. A reusable party identification block, for example, may combine established wording, alternatives for different party types, repetition for multiple parties, and a consistent presentation. Its reuse preserves the drafting decisions governing the text as well as the text itself.

At the architectural stage, the designer identifies the Snippet and the templates that need it. Its complete Content, Logic, and formatting are developed during Snippet Design. Where variable information must be inserted, Fields identify that information; where user input is needed, Questions are associated with the Snippet.

Sharing a Snippet allows the same legal expertise to be developed, reviewed, and maintained together across documents. Whether drafting can be shared depends on the legal function and required behaviour as well as the wording. Similar passages may require different treatment in different legal contexts, while different examples may reflect the same underlying drafting method. The comparison must therefore consider the Requirements of every template that will use the component.

Questions request the user input needed to generate documents. They identify information or choices that a user must supply for a Document Template or Snippet to operate as intended. Examples include a payment period, a choice between contractual alternatives, or confirmation of a circumstance that determines whether a provision applies.

Questions and Fields have distinct functions. Fields identify variable information to be inserted into the document from the Client Case. Questions request user input and connect that input to its use in the Document during detailed Design. Information already available in the Client Case does not, merely because it is needed for generation, create a need for

a Question.

An answer may supply information for a Field or inform Logic that includes a provision, selects an alternative, performs a calculation, or repeats content. At the architectural stage, the concern is to identify the user input need, give the Question a consistent name, and establish which templates or Snippets use it. Its final wording and the precise rules governing the use of its answer belong to detailed Design.

Questions can serve Document Templates directly and can also serve their Snippets. If a payment clause is shared across several templates, a Question requesting its payment period may belong to the design of that Snippet. Recording the Question's association with the Snippet makes this dependency visible alongside any direct use by templates. It also allows the same user input need to be recognized consistently wherever it recurs.

Shared components extend the value of each design decision across the portfolio. Without them, equivalent drafting and information needs may be addressed separately in every template. Repeated design effort can then produce divergent wording, rules, formatting, or requests for the same information.

Document Template Architecture makes those common needs visible. Consistent names identify the same Snippet or Question wherever it is used. Mapping Snippets to template Objectives shows which documents rely on shared drafting. Mapping Questions to template Objectives and Snippets shows where user input is required and which components depend on it.

These relationships also support maintenance. A change to a shared Snippet requires consideration of the templates that use it. A change to a Question may affect both templates and

Snippets, with consequences for the documents incorporating those Snippets. The Architecture gives the designer a basis for assessing these effects against the relevant Requirements and Acceptance Criteria.

Framework

Document Template Architecture begins with the Requirements of the selected Document Templates and produces a common record of Snippets, Questions, and their uses. The following steps identify recurring needs, name the components consistently, and make their dependencies explicit before detailed Design.

Step 1 — Compare the Selected Templates' Requirements

Review the Legal Task, Input Scope, and Output Characteristics of the selected Document Templates together. Identify recurring drafting needs and the information needed to address them. Acceptance Criteria help clarify the legal circumstances and observable results that the eventual design must support.

The comparison should consider Content, Logic, and formatting together. Two documents may require the same clause wording but different conditions for its use. Conversely, a common drafting method may support several documents through defined alternatives. The aim is to identify where a shared component can satisfy the relevant Requirements consistently.

Also distinguish information available from the Client Case from input that must be requested from a user. This keeps Questions connected to actual user input needs and leaves variable information to be specified through Fields during detailed Design.

Requirements and components need not have a one-to-one relationship. A single Snippet may address several required outcomes, while one outcome may depend on several Snippets and Questions. Some Content remains specific to an individual template. The comparison establishes the appropriate opportunities for reuse without requiring every part of a document to become a shared component.

Step 2 — Identify and Name Snippets and Questions

Identify the recurring combinations of Content, Logic, and formatting that should be developed as Snippets. Give each Snippet a consistent, descriptive name that expresses its drafting purpose. Use that name wherever the component appears in the Architecture and subsequent Design records.

Identify the user input needed by the templates and their Snippets, and give each Question a consistent name. The same Question should be recognizable wherever the same input need recurs. Similar wording alone does not establish that two Questions can be shared: the information requested and its meaning must be appropriate for each use.

Naming components consistently prevents equivalent needs from being treated as unrelated design tasks. It also distinguishes components that serve different legal purposes even where their wording is similar. Resolve material differences in intended use before treating a Snippet or Question as shared, so that detailed Design can take account of all its dependencies.

Step 3 — Map Components to Their Uses

For each Snippet, list the Objectives of the Document Templates that incorporate it. The Objective established during Selection and carried into Specification provides the reference linking the component to the relevant template and its Requirements.

For each Question, list the template Objectives and Snippets that use it. Record direct template use and use within Snippets as applicable. A Question used within a shared Snippet need not be represented as a separate direct use by every template incorporating that Snippet; the Snippet-to-template mapping makes the wider dependency visible.

Review the mappings against the selected templates' Requirements to confirm that the identified components serve the intended legal results. Check that shared components can accommodate the needs of all their uses and that Questions needed by Snippets have been included. These checks retain the connection to Specification while keeping the deliverable focused on components and dependencies.

The resulting map shows which components require coordinated design and which documents may be affected by a later change. It gives Document Template Design and Snippet Design a common structure for developing the detailed drafting method.

Resulting Deliverable

The resulting deliverable is the Document Template Architecture. It is a common record mapping named Snippets and Questions to their uses across the selected templates. Each Snippet is followed by the template Objectives that use it. Each Question is followed by the template Objectives and Snippets that use it, as applicable.

[snippet]

[template objective]

[...]

[question]

[template objective]

[...]

[snippet]

[...]

Repeat the Snippet and Question entries as needed. List only the uses that apply to each component: a Question may serve templates directly, Snippets, or both. Keep component names consistent with the names used in subsequent Design records, and retain the template Objectives used in the Requirements and Acceptance Criteria records.

This record makes reuse explicit before detailed Design begins. It shows where drafting decisions must serve several documents and where user input supports a shared component. Reviewing a component's listed uses gives the designer the context needed to develop it consistently and to assess the effect of later changes.

Conclusion

Document Template Architecture establishes shared Snippets and Questions by comparing the selected templates' Requirements, naming recurring components consistently, and mapping them to their uses. It connects the legal results defined during Specification to a common structure for detailed Design, extending the value of the firm's drafting expertise

across the portfolio.

Chapter 10 develops Document Template Design and Snippet Design. Each template and Snippet has a Document: the source file containing the Content, formatting, and Logic to which automation will be applied. The accompanying Design record identifies its Fields and Questions; a Document Template also identifies the Snippets it incorporates. The Document specifies where those Snippets are inserted and when they apply. Together, the source file and record translate the Architecture into instructions for implementation, with the completed asset subsequently assessed against its Acceptance Criteria.

10 Document Template Design

Chapter 9 established Document Template Architecture by comparing the selected templates' Requirements, identifying shared Snippets and Questions, and mapping those components to their uses. Snippets capture recurring Content, Logic, and formatting. Questions identify the user input needed by the templates and their Snippets. The Architecture makes these dependencies explicit before detailed Design begins.

Document Template Design and Snippet Design now specify how the drafting method will operate. Each Document Template and Snippet has a Document: the source file containing the Content, formatting, and Logic to which automation will be applied. Fields identify variable information drawn from the Client Case, while Questions connect required user input to its use in the Document. Document Templates also identify the Snippets they incorporate.

These designs preserve the lawyer's drafting choices as they are translated into automation. The source Document keeps the wording and structure connected to the conditions, selections, calculations, and repetitions governing the text. A record accompanying that source file identifies its associated components. Together, the Document and record give the implementer the instructions needed to reproduce the intended drafting method.

Concept

Requirements establish the Legal Task, Input Scope, and Output Characteristics. Acceptance Criteria define the Test Conditions and Expected Results used to evaluate the completed implementation. Architecture identifies shared components and their uses. Detailed Design specifies how the templates and Snippets will produce the required documents within that common structure.

A legal document combines several forms of drafting knowledge. Some wording is predetermined, some information varies with the matter, and some provisions depend on particular circumstances. Formatting organizes the resulting text, while rules govern which alternatives appear and how calculations and repetitions operate. Making this knowledge explicit allows established drafting decisions to be applied consistently across matters.

Document Automation applies Content and Logic through predominantly deterministic behaviour. A provision appears when a defined condition is satisfied, a controlled alternative is selected according to specified information, or a section repeats for each relevant party. The lawyer defines these relationships during Design so the implementer can reproduce them without reconstructing the underlying legal reasoning.

The Document is the source file containing the Content, formatting, and Logic to which automation will be applied. Each template and each Snippet has its own Document. It preserves the intended wording and presentation while specifying how the text must respond to the matter's information.

Content supplies the predetermined text: headings, introductory language, contractual provisions, definitions, notices, signature wording, and other controlled passages. Where approved wording embodies the firm's legal expertise, the

Document preserves that wording for repeated use. A provision can have predetermined Content even when its inclusion depends on Logic.

Formatting specifies the presentation of that Content. It can include heading styles, paragraph layout, numbering, lists, tables, emphasis, and the presentation of repeated sections. These choices form part of the intended document and must remain appropriate when optional text is omitted, alternatives are selected, or sections are repeated.

Logic defines the rules governing inclusion, selection, calculation, and repetition. Keeping those rules connected to the affected text makes the drafting method understandable in context. The source file should make clear which passage a condition governs, which alternatives belong together, and what structure is repeated. For a Document Template, it also specifies where each Snippet is inserted and when it applies.

Fields identify variable information drawn from the Client Case for insertion into the Document. Typical examples include party names, addresses, dates, references, amounts, and identifiers. Identifying a Field connects a location in the Document to the matter information required there. The design should make the intended information and its use unambiguous, including any necessary presentation such as a date format or currency representation.

Questions request user input needed to generate the document. They can supply information for a Field, determine whether a provision applies, select an alternative, or provide an input for a calculation. Their purpose extends to any required user input; it is not limited to decisions about document structure.

The design connects each Question to the place where its answer is used. It specifies what is requested, the permitted

form of the answer, and the consequence for the Document. Information already reliably available in the Client Case can be reused, avoiding unnecessary requests for the same information. Rules deriving a value from available inputs belong to Logic in the Document.

Document Template Design and Snippet Design use the same core elements: a Document, Fields, and Questions. A Snippet's Document contains its reusable Content, formatting, and Logic, while its associated Fields and Questions identify the variable information and user input it needs. A shared clause can therefore retain both its controlled wording and the rules governing its behaviour wherever it is incorporated.

Document Template Design additionally identifies the Snippets incorporated into the template. The template's Document specifies their locations and the circumstances in which they apply. The Snippet's own design specifies its internal drafting method. This division keeps shared decisions together while allowing the template to define how each component contributes to the complete document.

The Architecture supplies the common names and dependencies used in both designs. A shared component must be developed with all its uses in view. If detailed Design reveals a need to change a component or its uses, the Architecture and affected Design records should remain consistent so that the same name continues to identify the same component.

Framework

Document Template Design and Snippet Design develop the source Document and its associated components together. Begin with the Requirements, Acceptance Criteria, and Architecture, then work through the Document, Fields, Questions, and, for templates, incorporated Snippets. The design is com-

plete when these elements provide coherent instructions for implementation.

Step 1 — Develop the Document

Prepare the source Document for each template and Snippet. Specify the required Content and formatting, preserving the intended legal wording and organizing it into the appropriate document structure. Place shared drafting in the Snippet identified by the Architecture and identify its insertion point in the template's Document.

Make the Logic explicit alongside the text it governs. Four forms typically recur in Document Automation: conditions, selections, calculations, and repetitions. Their treatment should be precise enough for an implementer to determine the required behaviour from the design.

Conditions determine whether something is included or excluded. Specify the condition and the affected passage, heading, or Snippet, including the result when the condition is not satisfied. Where a Question is relevant only in defined circumstances, make that dependency clear as well.

Selections determine which predetermined alternative is used. Identify the alternatives and the circumstances selecting each one. The Document should preserve the wording of the alternatives and make their relationship clear so they can be implemented consistently.

Calculations derive a value from other information. Specify the inputs, formula, and any necessary treatment of units, rounding, or dates. Amounts, percentages, deadlines, periods, totals, and instalments can be calculated where the governing rule is defined in advance.

Repetitions apply a passage or structure to a collection of matter information. Specify what is repeated and which col-

lection supplies the entries, such as parties, items, transactions, or claims. The associated formatting should remain appropriate as the number of entries changes.

These forms of Logic can operate together. A repeated party identification block may select wording according to party type and include additional text only under a defined condition. Specify each rule at the level where it applies, keeping the inputs, affected text, and required result clear in the source Document.

Step 2 — Identify the Fields

Identify the Fields used in each Document and list them in its Design record. For each Field, make clear what variable information is drawn from the Client Case and where it is inserted. Use consistent names so repeated uses of the same information are recognizable.

Check the Fields against the Input Scope and the text requiring them. A reference to a party name, for example, must make clear which party is intended. Where a Field occurs within a repeated structure, its use must correspond to the relevant entry in the collection. Specify the presentation required for values such as dates and amounts in the Document.

Where user input supplies required information, connect that need to the relevant Question. Where a value is calculated, specify the calculation in the Document's Logic. These connections give the implementer a clear route from available information to its use in the generated text.

Step 3 — Specify the Questions

List the Questions needed by each template or Snippet using the names established in the Architecture. Specify the information requested and the permitted answer form, such as a

selection, date, amount, or text. Identify how each answer is used in the Document.

An answer may populate a Field, activate a Snippet, select a provision, determine whether another Question is relevant, or supply a calculation input. Make the corresponding rule explicit in the Document so that the request for information and its drafting consequence remain connected.

For Questions used by shared Snippets, consider the needs of all templates incorporating those Snippets. The Question and its use must remain consistent with the component's purpose. This preserves the benefit of shared information needs and prevents different templates from silently assigning different meanings to the same Question.

Step 4 — Identify the Snippets

For each Document Template, list the incorporated Snippets in its Design record. Check that the Document specifies where each is inserted and when it applies, using the same names as the Architecture and the corresponding Snippet Design records.

Review each Snippet's Document, Fields, and Questions alongside the templates that use it. Shared wording, Logic, and formatting must support the relevant Requirements in every use. Check that formatting, numbering, and references remain consistent as the document's content changes.

Resulting Deliverables

The deliverables are Document Template Design and Snippet Design. Each brings the source Document and its associated components together in a record accompanying the source file. The Document contains the Content, formatting, and Logic; the record identifies that Document and lists the Fields,

Questions, and, for a template, Snippets.

A Document Template Design record identifies the template by its Objective, preserving the connection to the selected Opportunity, Requirements, and Acceptance Criteria. It identifies the source Document and the Fields, Questions, and Snippets used by that template.

Document Template Design

Document Template — [objective]

Document

[Content, formatting and Logic]

Fields

[field]

[...]

Questions

[question]

[...]

Snippets

[snippet]

[...]

The Snippets listed in the record are developed through their own Snippet Design records. Their insertion locations and applicability are specified in the template's Document, keeping the assembly of the complete document connected to its drafting rules.

A Snippet Design record uses the component name established in the Architecture. It identifies the Snippet's own source Document, Fields, and Questions. Its Content, format-

ting, and Logic are specified in that source file, allowing the drafting method to be maintained consistently wherever the Snippet is used.

Snippet Design

[snippet]

Document

[Content, formatting and Logic]

Fields

[field]

[...]

Questions

[question]

[...]

The records can be supplemented with the detail needed to clarify Fields and Questions, while the source Documents preserve the actual wording, formatting, and drafting rules. The implementer should be able to follow their connections without reopening the legal analysis to decide what a provision should say, where information comes from, or how alternative outcomes are selected.

Conclusion

Document Template Design and Snippet Design complete the design of reusable drafting expertise. The Architecture identifies shared components and dependencies; the source Documents and their accompanying records make the operation of those components explicit. Implementation can then apply the specified Content, Logic, formatting, Fields, Ques-

tions, and Snippets to generate matter-specific documents.

After implementation, Acceptance Criteria establish whether the Requirements are met. Business Validation subsequently evaluates the business value realized through use of the asset. These connections retain the V-model's link between the implemented drafting method, its intended legal results, and the reasons for developing it.

Part V turns to Legal Intelligence Design for AI Workflows. Chapter 11 develops AI Workflow Architecture by identifying reusable Roles and Instructions and mapping them to the Workflows that use them. AI Workflow Design then organizes those components into connected Steps, while Role Design and Instruction Design specify their operation.

Part V

Legal Intelligence Design

11 AI Workflow Architecture

Selection identified which Opportunities should proceed to Specification. Requirements Determination then defined each selected Legal Asset's Legal Task, Input Scope, and Output Characteristics, while Acceptance Criteria Setting established the Test Conditions and Expected Results against which its implementation will be assessed. Together, these decisions establish what each asset must accomplish and how satisfactory performance will be evaluated.

Part V develops Legal Intelligence Design. For AI Workflows, Design begins with AI Workflow Architecture: identifying reusable Roles and Instructions and mapping them to the Workflows that use them. AI Workflow Design then organizes those components into connected Steps, while Role Design and Instruction Design specify their operation. Architecture provides the shared structure through which analytical approaches and assignments can be developed consistently across Workflows.

This connection matters for productization. An AI Workflow captures a legal analytical method that can be applied repeatedly across matters. Where several Workflows require the same analytical approach or assignment, shared components allow that expertise to be developed and maintained together. AI Workflow Architecture makes those opportunities for reuse and their dependencies explicit before detailed Design begins.

Concept

Requirements describe the intended legal result while leaving the mechanism for producing it open. The Legal Task establishes the work to be performed, the Input Scope identifies the necessary matter information and reusable legal knowledge, and the Output Characteristics describe what a satisfactory result must contain or demonstrate. Acceptance Criteria provide the corresponding basis for evaluation after implementation, allowing variation in wording while requiring the substantive qualities defined for the Workflow.

AI Workflow Architecture compares the selected Workflows' Requirements to identify the analytical approaches and assignments they need. Roles capture reusable analytical approaches, while Instructions specify reusable assignments. Mapping these components to the Workflows that use them shows which components should be developed together and which assets depend on the same expertise.

A Role defines the approach associated with a Legal Function: the professional activity performed by the model. Examples include identifying legally relevant issues, evaluating arguments, reviewing compliance, or synthesizing legal information. The Role establishes how the model should approach its work, giving different assignments a common analytical basis where appropriate.

At the architectural stage, the designer identifies the Role and the Workflows that need it. Role Design subsequently specifies its Legal Function, Reasoning Perspective, and Behavioural Constraints. These elements define the activity, the approach to interpreting and evaluating information, and the boundaries within which the work is performed.

An Instruction describes a Legal Assignment: the legal work to be performed. It gives substantive direction to the model's

work, such as identifying relevant contractual obligations, assessing them against a legal framework, or preparing a structured account of the findings. The Instruction establishes what the assignment must accomplish and how its result should be expressed.

At the architectural stage, the designer identifies the Instruction and the Workflows that need it. Instruction Design subsequently specifies its Legal Assignment, Applicable Rules, and Required Output. These elements define the work, the rules guiding its performance, and the result it must produce.

Separating Roles and Instructions supports reuse in two directions. The same analytical approach may support different assignments, and the same assignment may be needed in several Workflows. A Role used to review compliance, for example, may support different Instructions for examining different legal obligations. An Instruction for identifying contractual obligations may contribute to both a compliance review and a broader contractual risk analysis, where its assignment and required result suit both uses.

Reuse depends on the substance of the work. Similar names or broadly related topics do not establish that two Workflows need the same component. A Role must provide an appropriate analytical approach for every use, and an Instruction must describe an assignment that meets each Workflow's needs. Material differences in legal context, analytical purpose, or expected result may require distinct components.

This shared structure supports the broader objective of transforming legal expertise into reusable Legal Assets. Roles and Instructions are components of AI Workflows, one of the two kinds of Legal Asset addressed by Legal AI Design. Capturing expertise in named components allows improvements to an analytical approach or assignment to benefit every Workflow

that uses it.

The Architecture remains grounded in Specification. Each component must contribute to the legal work required by the Workflows to which it is mapped. A single component may support several required outcomes, while one outcome may depend on several components. Reviewing the mappings against the Requirements makes omissions and unsuitable reuse visible without requiring a separate component for every Requirement.

AI Workflow Design develops the execution structure from this common set of components. It distinguishes changes in assignment, analytical function, or required output through numbered Steps, each applying a Role and an Instruction from the Architecture. It also connects the results of those Steps: an intermediate result may inform a later analysis, and several results may contribute to the final output. This preserves the value of decomposing complex legal work into defined units while placing their organization in detailed Design.

Knowledge Cases are also identified for each Step during AI Workflow Design. They supply reusable legal knowledge such as legislation, case law, precedents, and internal guidance. Recording them at Step level allows the same Role and Instruction to be used with different knowledge sources. Matter information is supplied during execution. The Architecture maps the reusable analytical components to their Workflow uses; the subsequent Design records establish their operation and supporting knowledge.

Framework

AI Workflow Architecture begins with the Requirements of the selected AI Workflows and produces a common record

of Roles, Instructions, and their uses. The following steps identify recurring needs, name the components consistently, and make their dependencies explicit before detailed Design.

Step 1 — Compare the Selected Workflows' Requirements

Review the Legal Task, Input Scope, and Output Characteristics of the selected AI Workflows together. Identify the professional activities and legal assignments needed to produce the intended results. Acceptance Criteria help clarify the legal circumstances and substantive qualities that the eventual design must support.

Consider which analytical approaches recur and which assignments are required by more than one Workflow. Identifying relevant facts, applying a legal framework, evaluating an issue, and preparing findings may each reveal reusable expertise. The comparison should establish where that expertise can serve several Workflows consistently.

Assess the legal context and expected result as well as the general activity. Two Workflows may both involve reviewing contracts while requiring different analytical perspectives or different assignments. Conversely, a common assignment may serve Workflows with different overall Objectives. The comparison establishes where sharing a Role or Instruction is appropriate and where the legal work requires a distinction.

The Requirements remain the basis for these decisions. Components should have a clear contribution to the intended legal results. A component may be needed by only one selected Workflow, while others may already have several uses; identifying the common structure does not require every component to be shared immediately.

Step 2 — Identify and Name Roles and Instructions

Identify the Roles needed to provide the analytical approaches revealed by the comparison. Give each Role a consistent, descriptive name that expresses its professional activity. Identify the Instructions needed to capture the legal assignments, and give each a name that makes its assignment recognizable.

Consider existing components before introducing new ones. Where an existing Role provides the required analytical approach, reference it consistently. Where an existing Instruction captures an assignment suited to the Workflow's needs, use that Instruction. Review the component's intended operation to establish whether reuse is appropriate.

Resolve material differences before treating components as shared. Broadly similar activities may require different Reasoning Perspectives or Behavioural Constraints, while similarly named assignments may require different rules or outputs. These differences inform the decision to share a component, with its complete content subsequently specified through Role Design or Instruction Design.

Keep Roles and Instructions distinct when naming them. A Role should be recognizable as an analytical approach that can support appropriate assignments. An Instruction should be recognizable as the work to be performed. Consistent names allow the same component to be referenced wherever it appears in the Architecture and subsequent Design records.

Step 3 — Map Components to Their Workflow Uses

For each Role, list the Objectives of the AI Workflows that use it. For each Instruction, list the Workflow Objectives that use it. The Objective established during Selection and carried into Specification provides the reference linking each component to the relevant Workflow and its Requirements.

Review these mappings against the selected Workflows' Requirements. Confirm that the identified Roles and Instructions collectively address the analytical approaches and assignments needed for the Legal Task, taking account of the Input Scope and Output Characteristics. Check that each shared component can serve all its listed uses consistently.

The mappings provide traceability through the Workflow Objective without adding Step purposes, Requirement identifiers, or sequencing fields to the Architecture record. Detailed Design uses the identified components to organize the work into connected Steps and to specify how those Steps will produce the required results.

The resulting map also supports maintenance. A change to a Role or Instruction requires consideration of every Workflow that uses it. Its listed uses give the designer a basis for reviewing the consequences against the relevant Requirements and Acceptance Criteria, so that improvements remain appropriate across the portfolio.

Resulting Deliverable

The resulting deliverable is the AI Workflow Architecture. It is a common record mapping named Roles and Instructions to their uses across the selected Workflows. Each Role is followed by the Workflow Objectives that use it. Each Instruction is followed by the Workflow Objectives that use it.

[role]

[workflow objective]

[...]

[...]

[instruction]

[workflow objective]

[...]

[...]

Repeat the Role and Instruction entries as needed. List each component's applicable Workflow uses, keeping names consistent with the names used in subsequent Design records. Retain the Workflow Objectives used in the Requirements and Acceptance Criteria records so that each dependency can be traced to its legal purpose.

This record makes reuse explicit before detailed Design begins. It shows where an analytical approach or assignment must serve several Workflows and which components require coordinated development. Reviewing a component's listed uses gives the designer the context needed to develop it consistently and to assess the effect of later changes.

Conclusion

AI Workflow Architecture establishes reusable Roles and Instructions by comparing the selected Workflows' Requirements, naming the components consistently, and mapping them to the Workflow Objectives that use them. It connects the legal results defined during Specification to a shared struc-

ture for detailed Design, extending the value of the firm's analytical expertise across the portfolio.

Chapter 12 develops AI Workflow Design, Role Design, and Instruction Design. AI Workflow Design organizes the components into connected, numbered Steps and records the Role, Instruction, and Knowledge Cases for each Step under the Workflow's overall Objective. Role Design specifies Legal Function, Reasoning Perspective, and Behavioural Constraints. Instruction Design specifies Legal Assignment, Applicable Rules, and Required Output. Together, these records turn the shared Architecture into an implementation-ready design, with the implemented Workflow subsequently assessed against its Acceptance Criteria and its realized business value evaluated through Business Validation.

12 AI Workflow Design

Chapter 11 established AI Workflow Architecture by identifying reusable Roles and Instructions and mapping them to the Workflows that use them. Roles capture analytical approaches, while Instructions capture legal assignments. The Architecture makes their shared uses explicit and provides a common structure for detailed Design.

AI Workflow Design now organizes those components into connected, numbered Steps. Each Step applies a Role and an Instruction from the Architecture and identifies the Knowledge Cases needed to support its work. Role Design and Instruction Design specify the analytical approach and assignment in separate records, allowing both to be reused across Steps and Workflows.

Together, these designs translate the firm's analytical method into instructions for implementation. The Workflow record establishes how the work is divided and connected; the Role and Instruction records specify how each part of that work should be performed. Making these decisions explicit allows legal expertise to be applied repeatedly across matters while preserving the approach, rules, and output expectations that give it value.

Concept

Requirements establish the Legal Task, Input Scope, and Output Characteristics. Acceptance Criteria define the Test Con-

ditions and Expected Results used to evaluate the completed implementation. Architecture identifies reusable Roles and Instructions and their Workflow uses. Detailed Design specifies how those components will work together to produce the intended legal results.

An AI Workflow captures an analytical method for work that must respond to the particular matter. Large Language Models apply that method with probabilistic behaviour, so conclusions and wording may vary with the information supplied. Design gives the analysis a defined purpose and structure, while Acceptance Criteria establish the substantive qualities an acceptable result must demonstrate.

A Step is an individual unit of Workflow execution. Separating complex work into Steps gives each assignment the analytical approach, instructions, and knowledge it needs. A Step may identify relevant facts, apply a legal framework, evaluate an issue, prepare an intermediate analysis, or synthesize findings. Changes in assignment, analytical function, or required output provide the basis for distinguishing Steps.

AI Workflow Design determines the sequence of those Steps and connects their results. An earlier Step may identify information needed by a later analysis; several intermediate results may contribute to a final assessment. These connections make the progression of the legal work explicit and allow each Step's Required Output to support the work that follows.

The Role establishes how the model approaches its work. Role Design specifies three elements: Legal Function, Reasoning Perspective, and Behavioural Constraints. The Legal Function defines the professional activity performed by the model. The Reasoning Perspective specifies how it should interpret and evaluate information. Behavioural Constraints

establish the boundaries within which that activity is performed. Together, they give every Step using the Role a common analytical basis.

The Instruction establishes what the assignment must accomplish and how its result should be expressed. Instruction Design specifies three elements: Legal Assignment, Applicable Rules, and Required Output. The Legal Assignment describes the legal work to be performed, Applicable Rules guide its performance, and Required Output specifies the result to be produced.

The Workflow's overall Objective and an Instruction's Legal Assignment operate at different levels. The Objective identifies the intended result of the reusable AI Workflow and preserves its connection to Selection and Specification. A Legal Assignment describes the work performed whenever a Step applies that Instruction. Several assignments can contribute to one Workflow Objective, and the same assignment can support different Workflows where their needs correspond.

Knowledge Cases identify the reusable legal knowledge needed by each Step, such as legislation, case law, precedents, or internal guidance. They are recorded at Step level, separately from the Role and Instruction. This allows the same analytical approach and assignment to be used with different knowledge sources where appropriate. A Step that needs no supporting Knowledge Cases records that explicitly.

Matter information is supplied during execution from the Client Case. It may include facts, documents, correspondence, structured matter data, and previous work product within the Input Scope. The design must take account of what each assignment needs, while the reusable records retain

the distinction between the method and the information for a particular matter. Outputs developed during the Workflow can also inform subsequent Steps.

Legal Intelligence supports this contextual work using information from Client Cases and Knowledge Cases. A clear design identifies the supporting knowledge and the way each assignment uses its inputs, so the implementation can make the relevant information available. The method should not depend on an assumption that the model will independently possess all necessary facts or legal sources.

Separate Workflow, Role, and Instruction records preserve reuse. Each Step references its named Role and Instruction; their detailed designs are maintained in their own records. An improvement to a shared approach or assignment can therefore benefit every Workflow using it. The designer must consider all uses mapped in the Architecture when changing a shared component.

Framework

AI Workflow Design, Role Design, and Instruction Design develop the execution structure and its reusable components together. Begin with the Requirements, Acceptance Criteria, and Architecture. Organize the work into Steps, assign the components and supporting Knowledge Cases, and specify the Roles and Instructions until the records provide coherent instructions for implementation.

Step 1 — Organize and Connect the Workflow Steps

Retain the Workflow Objective established during Selection and used in its Requirements and Acceptance Criteria records. Review the Legal Task, Input Scope, and Output Characteristics to determine what work must occur to achieve that Objective.

Separate the work where the assignment, analytical function, or required output changes. Number the Steps in the order in which the work should proceed. A single Step may contribute to several required outcomes, while one required outcome may depend on several Steps. The decomposition should follow the structure of the legal work.

Determine how each result contributes to the subsequent work. A factual extraction may supply information for a legal assessment, and that assessment may support a final account of findings. Where a Step combines earlier results, make that use clear in its Instruction. Specify intermediate results through the relevant Instruction's Required Output so that each connection has a usable result.

Check that every Step contributes to the Workflow's Requirements and that the sequence collectively addresses the intended result. This gives each unit a defined purpose and retains traceability to Specification.

Step 2 — Assign Roles, Instructions, and Knowledge Cases

For each Step, identify the Role and Instruction from the Architecture, using the same component names. Check that the Role's analytical approach suits the Instruction's assignment and that the combination serves the Step's contribution to the Workflow.

Identify the Knowledge Cases needed to perform that assignment. Select the supporting legal knowledge in light of the Input Scope and the work required, and record the relevant Knowledge Cases or state that none are needed. Keep these references at Step level so that the Role and Instruction remain reusable with appropriate knowledge in other contexts.

Consider the matter information supplied during execution and any earlier Step results used by the assignment. Make

their intended use clear when specifying the Instruction, taking account of the Required Output needed by subsequent Steps. This connects the reusable method to its execution without incorporating the facts of an individual matter into the component designs.

Step 3 – Design the Roles

Develop a Role Design record for each named Role. Consider all its Workflow uses in the Architecture so the intended analytical approach can be applied consistently wherever that Role is assigned.

Specify the Legal Function: the professional activity performed by the model, such as identifying issues, evaluating arguments, reviewing compliance, or synthesizing legal information. Describe the activity precisely enough to orient the analysis while keeping the particular assignment in the Instruction.

Specify the Reasoning Perspective: how the model should interpret and evaluate information when performing that function. This may establish the standpoint for assessing relevance, weighing competing interpretations, recognizing uncertainty, or examining a legal position. The perspective should support the professional activity and remain appropriate across the Role's intended uses.

Specify the Behavioural Constraints: the boundaries governing that work. They may restrict assumptions, require gaps or uncertainty to be identified, limit the scope of conclusions, or establish how unsupported propositions must be treated. Express the intended behaviour clearly so the implementer can reproduce the lawyer's analytical approach.

Together, Legal Function, Reasoning Perspective, and Behavioural Constraints define the approach applied by each

Step using the Role. Review their consistency with the Instructions that the Role will support.

Step 4 — Design the Instructions

Develop an Instruction Design record for each named Instruction. Consider all its Workflow uses in the Architecture so that its assignment, rules, and expected result suit every use.

Specify the Legal Assignment: the legal work to be performed. Describe the assignment with enough precision to direct the model's work, including how it should use the relevant matter information or previous results. The assignment should contribute clearly to the Workflow's Legal Task and be compatible with the Role applied by the Step.

Specify the Applicable Rules: the legal, professional, methodological, or Workflow rules guiding performance of the assignment. They may establish the analytical sequence, relevant standards, treatment of conflicting information, or conditions governing conclusions. Keep the rules distinct from the supporting Knowledge Cases, which are identified separately for each Step.

Specify the Required Output: the nature and structure of the result to be produced. It may be intended for a lawyer, contribute to another Step, or constitute the Workflow's final result. Define the substantive content and presentation needed to make that result usable. Where later analysis depends on an intermediate result, ensure the Required Output provides what that analysis needs.

Required Output describes the result to be produced during execution. It is part of the reusable Instruction design; the completed matter-specific analysis is produced when the Workflow runs. Its design must allow the substantive quali-

ties required by Specification to be evaluated while accommodating the variation inherent in language-based analysis.

Resulting Deliverables

The deliverables are AI Workflow Design, Role Design, and Instruction Design. Together, these records describe the execution structure, analytical approaches, and assignments needed to implement the reusable Workflow. Component names remain consistent with the Architecture so that each reference identifies the corresponding detailed design.

AI Workflow Design

The AI Workflow Design record identifies the Workflow by its overall Objective. Under each numbered Step, it records the Role, Instruction, and Knowledge Cases, or states that no Knowledge Cases are needed. The record shows how the method and its supporting knowledge contribute to the intended result.

AI Workflow — [objective]

Step 1

[role]

[instruction]

[knowledge cases or none]

Step 2

[role]

[instruction]

[knowledge cases or none]

[...]

Repeat the Step entries as needed, retaining their execution order. Use the names of the Roles and Instructions identified in the Architecture and specified in the corresponding Design records. The Step's Knowledge Cases identify its supporting reusable legal knowledge; matter information is supplied during execution.

Role Design

Each Role Design record identifies a named Role and specifies its Legal Function, Reasoning Perspective, and Behavioural Constraints. It provides the common analytical basis for every Step that applies that Role.

[role]

Legal Function

[description]

Reasoning Perspective

[description]

Behavioural Constraints

[description]

Maintain a separate record for each Role. Review its content with all its mapped Workflow uses in view so that changes remain appropriate for the assignments it supports.

Instruction Design

Each Instruction Design record identifies a named Instruction and specifies its Legal Assignment, Applicable Rules, and Required Output. It makes the assignment reusable while leaving the choice of supporting Knowledge Cases with the individual Workflow Step.

Instruction Design

[instruction]

Legal Assignment

[description]

Applicable Rules

[description]

Required Output

[description]

Maintain a separate record for each Instruction. Its Required Output must support the use made of that result in every Workflow applying it. Where assignments require materially different rules or results, keep the component designs and Architecture mappings consistent with those differences.

Together, the records should allow an implementation specialist to configure the Workflow without reconstructing the underlying legal method. Technical decisions about models, retrieval, orchestration, interfaces, and integrations remain part of implementation. The design supplies the legal approach, assignments, supporting knowledge references, and output expectations those decisions must realize.

Conclusion

AI Workflow Design completes the execution structure by organizing reusable Roles and Instructions into connected, numbered Steps and identifying each Step's supporting Knowledge Cases. Role Design specifies Legal Function, Reasoning Perspective, and Behavioural Constraints. Instruction Design specifies Legal Assignment, Applicable Rules, and Required Output. These separate records make both the analytical approach and assignment reusable across matters and Workflows.

This completes the Design progression for AI Workflows. Selection identified the Opportunities worth pursuing; Specification defined their Requirements and Acceptance Criteria; Architecture identified reusable components and mapped their uses; detailed Design specified how those components perform the work. After implementation, Acceptance Criteria establish whether the Requirements are met, while Business Validation evaluates the business value realized through use.

The Conclusion brings these phases together as the Legal AI Design methodology. Document Templates preserve established drafting decisions in Content and Logic, while AI Workflows capture analytical methods that respond to each matter. Together, these Legal Assets allow the firm's expertise to be applied repeatedly, maintained as an institutional resource, and improved over time.

Conclusion

The promise of artificial intelligence in legal practice is often framed as a question of technology. Which model should a law firm use? Which tasks can be automated? Which platform should it buy? These questions matter, but they are not the central challenge. Technology does not determine which legal expertise should be captured, what a reusable legal solution must accomplish, how its quality should be evaluated, or how it should be structured for reliable repeated use. Those are design questions.

Legal AI Design addresses that design problem. It provides a methodology for transforming professional expertise that ordinarily resides in lawyers, documents, cases, and working practices into reusable legal assets. The objective is not automation for its own sake. It is the systematic productization of legal expertise.

Legal AI Design

Legal work contains substantial amounts of expertise that lawyers apply repeatedly. They identify relevant facts, select appropriate legal positions, determine which information matters, structure documents, formulate arguments, assess risks, and communicate conclusions. Much of this expertise remains implicit. It is embedded in experience, precedent documents, personal working methods, and the accumulated practices of individual lawyers and teams.

Legal AI Design provides a disciplined way to make that expertise explicit enough to be implemented as reusable legal assets. Its concern is therefore broader than the configuration of a particular technology. It addresses the decisions required to determine which applications of legal expertise deserve to become reusable assets, what those assets must accomplish, how their performance can be assessed, and how they should be structured and designed.

The V-model developed in this book provides the conceptual framework for that discipline. It brings together Selection, Specification, and Design while preserving the relationships between them. Selection concerns the reasons for developing a reusable legal asset in the first place. Specification makes its expected behaviour and quality explicit. Design translates those expectations into an implementable structure. The corresponding relationship between Requirements and Acceptance Criteria ensures that what is designed can ultimately be evaluated against what was required.

Within this framework, Opportunities make potentially reusable applications of legal expertise visible. Their frequency of reuse determines their relative priority within each Process and provides the basis for selecting which proceed to Specification. Business Validation assesses the business value realized through use of the assets after implementation.

Requirements express what a selected legal asset must accomplish, while Acceptance Criteria establish how satisfaction of those requirements can be assessed. Architecture defines the overall structure of the asset and the responsibilities of its constituent parts. Design provides the detailed specification required to implement its behaviour. These elements serve different purposes, but together they create the discipline needed to move from an idea about automation to an

implementation-ready design for a reusable legal asset.

Defining these elements before implementation is important because legal technology projects can otherwise move too quickly from perceived opportunity to development. A lawyer recognizes a repetitive task, a technology appears capable of performing it, and development begins. The result may function technically while solving the wrong problem, automating work with insufficient value, or producing outputs whose quality was never adequately specified. Legal AI Design places the design of the legal solution before its implementation.

The methodology applies to both forms of reusable legal assets examined throughout this book: Document Templates and AI Workflows. For predominantly deterministic Document Templates, designers can specify with considerable precision how particular inputs should affect the resulting document. For probabilistic AI Workflows, they must define the substantive qualities an acceptable result should demonstrate while allowing variation in wording.

Legal AI Design provides a common design discipline while allowing the resulting assets to be designed and evaluated according to their different characteristics. This makes it possible to approach legal automation as a coherent field of professional design rather than as a collection of unrelated technologies.

Productizing Legal Expertise

Legal AI Design produces implementation-ready designs for reusable Document Templates and AI Workflows, which become operational Legal Assets through implementation. Together, they provide two complementary mechanisms for capturing legal expertise that would otherwise have to be

repeatedly reconstructed in individual matters.

Document Templates capture predominantly deterministic applications of legal expertise. A well-designed template contains more than reusable text. It embodies decisions about document structure, required information, alternative provisions, conditional content, calculations, repetitions, and the relationships between these elements. Expertise that previously depended on a lawyer knowing which precedent to select, which clauses to include, and how to adapt them can instead be represented through predefined content and logic.

AI Workflows extend this principle to applications of expertise that cannot readily be expressed through deterministic rules alone. Much legal work depends on interpreting language and context. Lawyers summarise documents, identify relevant facts, compare positions, classify information, analyse arguments, transform unstructured material, and produce drafts from complex source information. AI Workflows make it possible to capture the instructions, context, constraints, and expected outputs associated with these activities so that generative AI can assist in applying that expertise repeatedly.

The distinction between deterministic and probabilistic legal assets is therefore not merely technical. It defines two complementary ways of making professional expertise reusable. Document Templates are particularly suitable where the relationship between information and output can be specified in advance. AI Workflows are particularly suitable where the application of expertise depends on language, interpretation, or contextual reasoning. Many areas of legal work contain both.

This changes the location of legal expertise within the firm.

Traditionally, much of a firm's expertise is distributed. Part of it resides in the minds of experienced lawyers. Part of it exists in precedent documents. Part can be reconstructed from previous matters. Another part survives as informal knowledge about how particular tasks are usually performed. Even when that expertise is repeatedly applied, it does not necessarily exist as a reusable asset in its own right.

Productization changes this relationship. Expertise is deliberately captured in a form that can be repeatedly applied. A carefully designed Document Template does not merely preserve the output of previous legal work; it preserves relevant aspects of the logic required to produce future outputs. An AI Workflow does not merely preserve a useful prompt; it specifies a repeatable application of legal expertise within a defined context and for a defined purpose.

The consequence is cumulative. Each application of Legal AI Design can add another reusable asset to the firm's capabilities. Existing assets can be maintained as law, practice, client requirements, or technology change. Their Requirements and Acceptance Criteria provide a basis for assessing modifications. Their Architecture and Design make their internal structure explicit enough to support controlled improvement. New assets can be developed where additional valuable applications of legal expertise are identified.

Over time, the result is an expanding body of productized expertise. Instead of repeatedly rebuilding the same professional knowledge within individual matters, the firm increasingly applies, maintains, and improves assets in which that knowledge has already been captured.

This is the meaning of productizing legal expertise at the core of this book. Productization does not require turning every legal service into a standardized commodity, nor does it im-

ply removing lawyers from legal work. It means identifying professional expertise that is repeatedly applied and transforming appropriate parts of that expertise into reusable legal assets. Document Templates and AI Workflows make that transformation technically possible. Legal AI Design makes it systematic.

Commercial and Strategic Advantage

Reusable legal assets have consequences beyond operational automation. Once legal expertise can be repeatedly applied through Document Templates and AI Workflows, those assets can become the foundation for productized legal services.

The immediate benefits are consistency, efficiency, and predictability. Reusable assets can reduce unnecessary reconstruction of work that the firm has effectively performed before. They can help ensure that established knowledge and working methods are applied more consistently across matters. They can also make the resources required for recurring work more predictable because part of the professional effort has been transferred from repeated execution into the design and maintenance of the underlying asset.

This changes the economics of legal work. Traditional legal services often maintain a close relationship between professional time spent and revenue generated. Increasing efficiency can therefore create an uncomfortable commercial tension: performing the same work in less time may reduce the amount that can be billed when price remains directly tied to hours worked.

Productized legal services create the possibility of weakening that relationship. When a meaningful part of the value delivered to a client comes from reusable assets developed in advance, the professional time required for an individual

matter no longer represents the full productive input behind the service. The client benefits from accumulated expertise that can be applied efficiently and consistently, while the firm can explore pricing models that reflect the value of the service rather than only the time required to execute it.

For clients, this can mean more predictable service delivery, greater consistency, faster execution, and clearer relationships between needs, outputs, and price. Productization can make recurring legal work easier to understand and procure without requiring the underlying legal expertise to become simplistic.

For law firms, the significance is deeper. Expertise that exists only in individual lawyers is difficult to scale. Expertise embedded in reusable assets can be applied across lawyers, teams, matters, and clients. Professional knowledge begins to acquire some of the characteristics of organizational capital: it can be retained, deployed, maintained, and improved independently of the individual matter in which it was originally developed.

As this organizational capital accumulates, it becomes a source of strategic advantage. Most law firms can acquire access to similar foundation models, document automation systems, and other legal technologies, so technology itself provides a limited basis for sustainable differentiation. What competitors cannot acquire as easily is another firm's accumulated body of productized legal expertise.

A mature collection of Document Templates and AI Workflows reflects years of substantive knowledge, client experience, drafting choices, quality standards, working methods, and design decisions. Each asset may represent a relatively narrow application of expertise, but their combined value can become substantial. As the collection grows and im-

proves, the firm develops a proprietary capability that is increasingly embedded in reusable assets rather than remaining dependent entirely on individual execution.

The strategic question for law firms is therefore not simply how much artificial intelligence they use. It is how much of their valuable, repeatedly applied expertise they have succeeded in turning into assets.

The opportunity is to make that expertise scalable by systematically capturing it: combining professional judgment with reusable assets and building firm-specific capabilities that grow with each application of the methodology.

Technology can be purchased. Productized expertise must be built.

Legal AI Design provides the methodology for building it.

Glossary

Acceptance Criteria. Standards used to assess whether an implemented Legal Asset meets its Requirements. Each criterion pairs Test Conditions with an Expected Result.

Acceptance Criteria Setting. The activity within Specification that defines the evidence needed to accept an asset's performance, using criteria suited to its behaviour.

AI Workflow. A reusable Legal Asset that captures an analytical method in structured Steps. It uses Large Language Models to perform contextual analysis and language-based legal work using the information available for a matter.

AI Workflow Architecture. The shared structure identifying reusable Roles and Instructions and mapping them to the AI Workflows that use them.

AI Workflow Design. The detailed design that organizes an AI Workflow into numbered Steps and identifies the Role, Instruction, and supporting Knowledge Cases for each Step.

Annual Frequency. The number of times a Legal Asset is used during a year, recorded as part of Business Validation.

Annual Value. The economic value a Legal Asset generates through use during a year, calculated as Value per Use × Annual Frequency.

Applicable Rules. The rules specified in an Instruction that guide the performance of its Legal Assignment.

Architecture. The structure that identifies shared components and maps their dependencies before detailed Design. Document Template Architecture identifies Snippets and Questions; AI Workflow Architecture identifies Roles and Instructions.

Behavioural Constraints. The boundaries established in a Role for how the model performs its work.

Business Validation. The evaluation, after implementation, of the business value realized through a Legal Asset's use. It records Value per Use, Annual Frequency, and Annual Value, connecting results to the potential identified during Selection.

Case. A unit within Case Management that brings together either matter-specific information and work in a Client Case or reusable legal knowledge in a Knowledge Case.

Case Management. The part of a Legal AI System that organizes legal work and information in Client Cases and Knowledge Cases.

Client Case. A Case containing the information and work associated with an individual client matter. It supplies matter-specific information for Document Automation and Legal Intelligence.

Content. The predetermined text captured in a Document Template or Snippet, whose use is governed by Logic.

Design. The phase of Legal AI Design that specifies how Legal Assets will meet their Requirements and Acceptance Criteria through Architecture and detailed designs ready for implementation.

Deterministic Behaviour. Behaviour in which predefined Content and Logic determine the result from the supplied

information. Document Automation operates predominantly in this way.

Document. In Document Template Design and Snippet Design, the source file containing the Content, formatting, and Logic to which automation will be applied. It specifies how the wording and structure depend on conditions, selections, calculations, and repetitions.

Document Automation. The part of a Legal AI System that applies a Document Template's Content and Logic to Client Case information to generate a matter-specific document.

Document Automation Design. The design of Document Templates and their reusable components, covering Architecture and the detailed specification of Content, Logic, formatting, Fields, Questions, and Snippets.

Document Template. A reusable Legal Asset that captures established drafting decisions in predefined Content and Logic, enabling documents to be generated for the circumstances of individual matters.

Document Template Architecture. The shared structure identifying reusable Snippets and Questions and mapping them to the Document Templates and components that use them.

Document Template Design. The detailed specification of a Document Template's Document, Fields, Questions, and incorporated Snippets. The Document specifies where Snippets are inserted and when they apply.

Expected Result. The outcome an evaluator must observe under the Test Conditions of an acceptance criterion. For AI Workflows, it can specify substantive qualities while allowing variation in wording.

Field. An element identifying variable information drawn from a Client Case for insertion into a Document.

Formatting. The presentation and layout specified alongside Content and Logic in a Document Template or Snippet. It can also be assessed through Acceptance Criteria.

Input Scope. The matter information, legal knowledge, and other inputs required to complete the Legal Task defined in the Requirements.

Instruction. A reusable assignment applied within an AI Workflow Step. It specifies the Legal Assignment, Applicable Rules, and Required Output.

Instruction Design. The detailed specification of what an assignment must accomplish and how its result should be expressed, through its Legal Assignment, Applicable Rules, and Required Output.

Knowledge Case. A Case containing reusable legal knowledge for use across matters, such as legislation, case law, precedents, or internal guidance. Supporting Knowledge Cases are identified separately for each AI Workflow Step.

Large Language Model. An AI model that processes and generates language. AI Workflows use these models to apply structured analytical methods with probabilistic behaviour.

Legal AI Design. The methodology for transforming legal expertise into reusable Document Templates and AI Workflows through Selection, Specification, and Design, linked to evaluation through the V-model.

Legal AI System. The environment connecting Case Management, Document Automation, and Legal Intelligence so that matter information, resulting work, and reusable assets can support the application and refinement of legal expertise.

Legal Asset. A reusable embodiment of legal expertise. The two kinds addressed by Legal AI Design are Document Templates and AI Workflows.

Legal Assignment. The legal work an Instruction requires the model to perform.

Legal Domain. A practice area or legal niche used to group recurring Processes and identify Opportunities for production.

Legal Function. The professional activity the model performs, defined as part of a Role.

Legal Intelligence. The part of a Legal AI System that supports contextual analysis and language-based legal work using information from Client Cases and Knowledge Cases. It can also help develop reusable assets.

Legal Intelligence Design. The design of AI Workflows and their reusable analytical components, covering AI Workflow Architecture, AI Workflow Design, Role Design, and Instruction Design.

Legal Task. The work an asset must support and the circumstances in which it will be performed, defined in the Requirements by translating the selected Objective into a legal need.

Logic. The predefined rules governing the inclusion, selection, calculation, and repetition of Content in a Document Template or Snippet.

Objective. A concise statement of an Opportunity's intended purpose: "Generate [document type] for [legal context]" for a Document Template, or "[verb] [noun] for [legal context]" for an AI Workflow.

Opportunity. An identified prospect for capturing recurring

legal expertise in a Document Template or AI Workflow, expressed through an Objective and grouped by Legal Domain and Process.

Opportunity Identification. The activity within Selection that identifies recurring legal work suitable for reusable assets and prioritizes Opportunities by frequency of reuse within each Process.

Output Characteristics. The qualities that a satisfactory result must contain or demonstrate, specified in the Requirements to guide Design and provide a basis for Acceptance Criteria.

Precedent. Earlier wording preserved for reuse. A Document Template additionally captures the decisions governing how wording should change to suit a matter.

Probabilistic Behaviour. Behaviour in which conclusions and wording can vary as a model applies instructions to contextual information. AI Workflows exhibit this behaviour through their use of Large Language Models.

Process. A recurring sequence or category of legal work within a Legal Domain, used to group and prioritize Opportunities.

Productization. The transformation of legal expertise into reusable Legal Assets that allow it to be applied across matters with less recurring lawyer effort, supporting more standardized and scalable legal work.

Question. A component requesting user input needed to generate a document and connecting that input to its use in a Document Template or Snippet.

Reasoning Perspective. The element of a Role specifying how the model should interpret and evaluate information.

Required Output. The result an Instruction specifies must be produced, including how that result should be expressed.

Requirements. The specification of what a Legal Asset must accomplish, comprising its Legal Task, Input Scope, and Output Characteristics.

Requirements Determination. The activity within Specification that translates a selected Objective into the Legal Task, Input Scope, and Output Characteristics needed to guide Design and evaluation.

Role. A reusable analytical approach applied within an AI Workflow Step, defined by its Legal Function, Reasoning Perspective, and Behavioural Constraints.

Role Design. The detailed specification of a Role's Legal Function, Reasoning Perspective, and Behavioural Constraints so that its analytical approach can be applied consistently across assignments.

Selection. The phase of Legal AI Design that determines which recurring legal work warrants investment in production. Opportunity Identification establishes priorities; Business Validation later evaluates the value realized through use.

Snippet. A reusable drafting component containing Content, Logic, and formatting that can be incorporated into Document Templates.

Snippet Design. The detailed specification of a Snippet's Document, Fields, and Questions, preserving its drafting choices as instructions for implementation.

Specification. The phase of Legal AI Design that defines what selected assets must accomplish and how satisfactory performance will be assessed, through Requirements Deter-

mination and Acceptance Criteria Setting.

Step. An individual unit of AI Workflow execution that applies a Role and an Instruction. Steps distinguish changes in assignment, analytical function, or required output and identify any supporting Knowledge Cases.

Test Conditions. The representative inputs and legal circumstances under which an asset is assessed against an acceptance criterion.

V-model. The framework connecting definition and design decisions to validation. Its descending side moves from Opportunities through Requirements to Architecture and Design; after implementation, Acceptance Criteria assess fulfillment of Requirements and Business Validation evaluates realized business value.

Value per Use. The economic value realized each time a Legal Asset is used, based on attributable revenue or the cost reduction achieved compared with equivalent manual work.

About the Book

Legal expertise typically depends on lawyer time, limiting how much legal work can be delivered with the professionals available.

Legal AI Design shows how to capture that expertise in reusable Document Templates and AI Workflows that can be applied across matters with less recurring effort.

This can increase revenue or reduce costs while building a growing portfolio of the organization's own expertise as a lasting strategic advantage.

About the Author

Dr. Gert Loterman is the creator of Legal AI Design and founder of Innoverius. He is an engineer focused on legal technology and holds a PhD in applied economics with a specialization in artificial intelligence.