UNIVERSIDAD PONTIFICIA COMILLAS
ESCUELA TÉCNICA SUPERIOR DE INGENIERÍA (ICAI)
GRADO EN INGENIERÍA EN TECNOLOGÍAS DE TELECOMUNICACIÓN Y BUSINESS ANALYTICS

TRABAJO FIN DE GRADO

CONCOORD: AN AGENTIC CONTRACT-MANAGEMENT PLATFORM FOR RENEWABLE-ENERGY INFRASTRUCTURE

Autor: José Guardo Medina
Director: David Martín-Corral Calvo

Curso académico 2025/26

Madrid

DECLARACIÓN DE AUTORÍA

Declaro, bajo mi responsabilidad, que el Proyecto presentado con el título

ConCoord: An Agentic Contract-Management Platform for Renewable-Energy Infrastructure

en la ETS de Ingeniería – ICAI de la Universidad Pontificia Comillas en el curso académico 2025/26 es de mi autoría, original e inédito y no ha sido presentado con anterioridad a otros efectos.

El Proyecto no es plagio de otro, ni total ni parcialmente, y la información que ha sido tomada de otros documentos está debidamente referenciada.

Fdo.: José Guardo Medina   Fecha: 22 / 04 / 2026

Autorizada la entrega del proyecto

EL DIRECTOR DEL PROYECTO

Fdo.: David Martín-Corral Calvo   Fecha: ___ / ___ / ___

MARTIN CORRAL CALVO DAVID - 03909768K

Digitally signed by MARTIN CORRAL CALVO DAVID - 03909768K
Date: 2026.04.21 18:24:10 +02'00'

UNIVERSIDAD PONTIFICIA COMILLAS
ESCUELA TÉCNICA SUPERIOR DE INGENIERÍA (ICAI)
GRADO EN INGENIERÍA EN TECNOLOGÍAS DE TELECOMUNICACIÓN Y BUSINESS ANALYTICS

TRABAJO FIN DE GRADO

CONCOORD: AN AGENTIC CONTRACT-MANAGEMENT PLATFORM FOR RENEWABLE-ENERGY INFRASTRUCTURE

Autor: José Guardo Medina
Director: David Martín-Corral Calvo

Curso académico 2025/26

Madrid

CONCOORD: AN AGENTIC CONTRACT-MANAGEMENT PLATFORM FOR RENEWABLE-ENERGY INFRASTRUCTURE

Author: José Guardo Medina
Supervisor: David Martín-Corral Calvo
Collaborating Entity: ICAI – Universidad Pontificia Comillas

ABSTRACT

ConCoord is an agentic software platform that continuously links the physical, daily-changing state of infrastructure assets to their entire contractual framework. This TFG builds the complete commercial and technical case for deploying ConCoord in the Spanish renewable-energy market, where over 120 Independent Power Producers manage contract portfolios of 50–100+ documents per plant across construction, operations, finance, and offtake. The work combines market analysis, competitive positioning, product and architectural design, pricing strategy, and a five-year financial model, supported by a go-to-market plan anchored in the Madrid-centred IPP community.

Keywords: Agentic AI, Contract Management, Renewable Energy, LLM, Multi-agent Systems, Infrastructure.

1. Introduction

Every infrastructure asset operates under a dense web of contracts with hard economic and legal consequences, yet the link between physical reality and contractual obligations is managed manually and reactively — losing the sector millions annually to preventable disputes, missed claims, and compliance failures. ConCoord is the agentic layer that closes this gap in real time.

2. Definition of the project

The project defines a full commercial and technical proposal rather than a deployed product. It covers problem scoping, market sizing, competitive landscape, technological foundations (agentic AI, knowledge graphs, cloud infrastructure), a detailed architecture specification, three complementary pricing strategies, a five-year unit-economics model, a 90-day deployment plan, and an alignment with the United Nations Sustainable Development Goals.

3. Description of the system / tool

ConCoord is composed of two tightly integrated pipelines. A design-time pipeline ingests contracts, decomposes them into clauses, maps cross-references through a knowledge graph, and runs a twelve-agent analysis that produces an executable Blueprint — an annotated state machine modelling the full lifecycle of every clause. A runtime pipeline then instantiates each Blueprint against a signed contract, driving obligations through pending → in-progress → fulfilled states while a deterministic Reminder Scheduler handles every deadline.

4. Results

At the base-case projection, ConCoord reaches 230 active plants in Year 5, generating $11.0M ARR, 91% gross margin, and 80% EBITDA margin. A representative 8-plant IPP client captures approximately $2.16M/year of protected value against a $403K/year Strategy 2 subscription — a 5.4× ROI multiple. Three pricing strategies (per-plant, SaaS subscription with volume tiers, and consultancy-level managed service) cover the range of buyer profiles in the Spanish market.

5. Conclusions and insights

The combination of agentic AI and deterministic runtime mechanics allows ConCoord to deliver something no incumbent in the consultancy, legal, or CLM software segments delivers today: real-time, evidence-grounded, proactive contract compliance. The Spanish renewable-IPP market is the right beachhead thanks to its fragmentation, capital intensity, and contractual density; the same logic extends naturally into data centres, conventional energy, and transport infrastructure.

6. References

The TFG cites 34 sources including the FIDIC contract suite [20][21], EU Directive 2018/2001 [18], Spanish Royal Decree 960/2020 [19], OMIE/MIBEL market rules [6], and academic work on agentic AI and legal NLP [27]–[31]./

CONCOORD: UNA PLATAFORMA AGÉNTICA DE GESTIÓN DE CONTRATOS PARA INFRAESTRUCTURAS DE ENERGÍA RENOVABLE

Autor: José Guardo Medina
Director: David Martín-Corral Calvo
Entidad Colaboradora: ICAI – Universidad Pontificia Comillas

RESUMEN DEL PROYECTO

ConCoord es una plataforma de software agéntica que vincula de forma continua el estado físico, cambiante día a día, de los activos de infraestructura con su marco contractual completo. Este TFG construye el caso comercial y técnico completo para desplegar ConCoord en el mercado español de energía renovable, donde más de 120 productores independientes de energía (IPPs) gestionan carteras contractuales de 50–100+ documentos por planta entre construcción, operación, financiación y comercialización. El trabajo combina análisis de mercado, posicionamiento competitivo, diseño de producto y arquitectura, estrategia de precios y un modelo financiero a cinco años, respaldado por un plan de salida al mercado anclado en la comunidad IPP de Madrid.

Palabras clave: IA agéntica, gestión de contratos, energía renovable, LLM, sistemas multi-agente, infraestructura.

1. Introducción

Todo activo de infraestructura opera bajo una red densa de contratos con consecuencias económicas y legales exigentes, pero el enlace entre la realidad física y las obligaciones contractuales se gestiona de forma manual y reactiva — perdiendo millones anualmente en disputas evitables, reclamaciones no presentadas y fallos de cumplimiento. ConCoord es la capa agéntica que cierra esa brecha en tiempo real.

2. Definición del proyecto

El proyecto define una propuesta comercial y técnica completa, no un producto desplegado. Abarca el encuadre del problema, dimensionamiento del mercado, panorama competitivo, fundamentos tecnológicos (IA agéntica, grafos de conocimiento, infraestructura cloud), una especificación arquitectónica detallada, tres estrategias de precios complementarias, un modelo de unit economics a cinco años, un plan de despliegue de 90 días y la alineación con los Objetivos de Desarrollo Sostenible de las Naciones Unidas.

3. Descripción del sistema / herramienta

ConCoord se compone de dos pipelines estrechamente integrados. Un pipeline de diseño (design-time) ingesta los contratos, los descompone en cláusulas, mapea las referencias cruzadas mediante un grafo de conocimiento y ejecuta un análisis de doce agentes que produce un Blueprint ejecutable — una máquina de estados anotada que modela el ciclo de vida completo de cada cláusula. Un pipeline de ejecución (runtime) instancia después cada Blueprint contra un contrato firmado, haciendo avanzar las obligaciones por los estados pendiente → en curso → cumplida mientras un Programador de Recordatorios determinista gestiona cada plazo.

4. Resultados

En el escenario base, ConCoord alcanza 230 plantas activas en el Año 5, generando 11,0 M$ de ARR, 91% de margen bruto y 80% de margen EBITDA. Un cliente IPP representativo con 8 plantas captura aproximadamente 2,16 M$/año de valor protegido frente a un coste de suscripción Estrategia 2 de 403 K$/año — un múltiplo de ROI de 5,4×. Las tres estrategias de precios (por planta, suscripción SaaS con tramos por volumen y servicio gestionado estilo consultoría) cubren el rango de perfiles de comprador del mercado español.

5. Conclusiones

La combinación de IA agéntica y mecánica determinista en tiempo de ejecución permite a ConCoord entregar algo que ningún actor existente en los segmentos de consultoría, boutique legal o CLM hace hoy: cumplimiento contractual proactivo, en tiempo real y con trazabilidad evidencial. El mercado español de IPPs renovables es la cabeza de playa adecuada por su fragmentación, intensidad de capital y densidad contractual; la misma lógica se extiende de forma natural a centros de datos, energía convencional e infraestructura de transporte.

6. Referencias

El TFG cita 34 fuentes, incluyendo la familia contractual FIDIC [20][21], la Directiva UE 2018/2001 [18], el Real Decreto 960/2020 [19], las reglas del mercado OMIE/MIBEL [6] y trabajos académicos sobre IA agéntica y PLN jurídico [27]–[31].

Índice de la memoria

Índice de figuras

Índice de tablas

Índice de abreviaturas

Abreviatura	Definición
AI	Artificial Intelligence — Inteligencia Artificial
API	Application Programming Interface
ARR	Annual Recurring Revenue — Ingresos Recurrentes Anuales
AWS	Amazon Web Services
CAPEX	Capital Expenditure — Gasto de capital
CLM	Contract Lifecycle Management — Gestión del ciclo de vida contractual
COGS	Cost of Goods Sold — Coste de los bienes vendidos
CEO	Chief Executive Officer
CTO	Chief Technology Officer
DPA	Data Processing Agreement — Acuerdo de tratamiento de datos
EBITDA	Earnings Before Interest, Taxes, Depreciation and Amortization
EKS	Elastic Kubernetes Service (AWS)
ECS	Elastic Container Service (AWS)
EPC	Engineering, Procurement and Construction
FIDIC	Fédération Internationale des Ingénieurs-Conseils
GDPR	General Data Protection Regulation
IPP	Independent Power Producer — Productor Independiente de Energía
ISO	International Organization for Standardization
KPI	Key Performance Indicator — Indicador clave de desempeño
LLM	Large Language Model — Modelo de lenguaje de gran tamaño

Abreviatura	Definición
LOI	Letter of Intent — Carta de intenciones
MIBEL	Mercado Ibérico de la Electricidad
MSA	Master Services Agreement
MVP	Minimum Viable Product — Producto Mínimo Viable
NDA	Non-Disclosure Agreement — Acuerdo de confidencialidad
O&M	Operation and Maintenance — Operación y Mantenimiento
ODS	Objetivo de Desarrollo Sostenible
OMIE	Operador del Mercado Ibérico de Energía
OPEX	Operating Expenditure — Gasto operativo
PLN	Procesamiento de Lenguaje Natural — Natural Language Processing
PPA	Power Purchase Agreement — Contrato de compraventa de energía
RDS	Relational Database Service (AWS)
RAG	Retrieval-Augmented Generation
ROI	Return on Investment — Retorno de la inversión
S3	Simple Storage Service (AWS)
SaaS	Software as a Service
SES	Simple Email Service (AWS)
SLA	Service-Level Agreement
SNS	Simple Notification Service (AWS)
SOC 2	Service Organization Control 2
TFG	Trabajo Fin de Grado
UE	Unión Europea
WAF	Web Application Firewall

Chapter 1. INTRODUCTION

1.1 Context and motivation of the project

Every infrastructure asset — whether under construction or in operation — is governed by a dense web of contracts: construction, maintenance, lease, supplier, customer. Each of those contracts carries obligations with hard economic and legal consequences tied to specific milestones, deadlines, and performance thresholds.

Yet today, there is no live connection between what is actually happening on the ground and what the contracts require. Tracking compliance is manual, fragmented, and reactive, meaning misalignments are discovered after they have already caused damage.

The result is that the infrastructure sector loses millions annually to preventable delays, missed or wrongfully executed claims, misunderstood rights, and compliance failures. [3, 4, 5]

This Final Degree Project (TFG) presents ConCoord, an agentic layer that continuously links the physical, daily-changing state of any infrastructure asset to its entire contractual framework, surfacing risks, obligations, and opportunities in real time. The work documents the full commercial and technical case for deploying such a platform, starting in the Spanish renewable-energy market, where contract density, capital intensity, and regulatory complexity make the opportunity especially acute.

1.2 General approach and scope

The system closes the contract–reality gap end to end. It ingests contracts in their native formats, converts every substantive clause into an executable state machine — called a Blueprint — and then instantiates that Blueprint as a live monitoring process that tracks obligations, enforces deadlines, and maintains a complete audit trail throughout the contractual lifecycle. The platform is composed of two tightly integrated pipelines: a design-time pipeline that extracts the Blueprint from the clause text, and a runtime pipeline that drives a Blueprint through its lifecycle for a specific real-world contract.

The scope of the present TFG is not a deployed commercial product but a complete project proposal: market analysis, competitive positioning, product and architectural description, pricing and financial modelling, unit economics, and a go-to-market plan. The deliverable is therefore a decision-grade business case supported by quantitative modelling, rather than a software artefact.

1.3 Technological, industrial and regulatory justification

The contractual risk that ConCoord addresses is systemic and growing, and it is disproportionately concentrated in the energy sector. Recent industry evidence illustrates the scale of the problem:

Statistic	**Figure**	**Ref**
Average U.S. construction dispute value	$60.1M (+40% YoY)	[3]
Projects that result in a claim	1 in 4	[4]
Total claims across 2,200+ analysed projects ($2.43T capex)	$95 billion	[4]
Disputed sums as share of contract budget	33.4%	[4]
Avg dispute resolution time	12.5 months	[3]
Leading cause of disputes (3 consecutive years)	Errors & omissions in contract documents	[3]
Energy companies: average annual litigation spend	$41 million	[5]
Energy-sector premium vs. other sectors for $15M+/yr litigation	56% more likely	[5]
Energy companies with $50M+ unenforced judgments sitting idle	35%	[5]
EPC liquidated-damages cap (FIDIC market practice)	10% of contract price	[20]
Global CLM market size (2026) → projected (2036)	$1.8B → $5.4B	[26]

Technically, the maturation of large language models and multi-agent orchestration frameworks has created a new design space for this problem — one that did not exist when incumbent CLM vendors were built. [27, 28] Industrially, Spain is one of Europe's leading renewable-energy markets [14, 15, 16], with large-scale deployment across solar PV, onshore wind, and a rapid buildout of hybrid PV+BESS assets [17]. Regulatorily, the overlapping frameworks that govern the value chain — EU Directive 2018/2001 [18], Spanish Royal Decree 960/2020 [19], FIDIC model conditions [20, 21], and MIBEL/OMIE market rules [6] — multiply the surface area of obligations that any single asset operator must track. The intersection of these three vectors is the justification for building ConCoord now.

Chapter 2. DESCRIPTION OF THE TECHNOLOGIES

ConCoord is built on three complementary technology layers: (i) agentic artificial intelligence, which provides the reasoning capacity to interpret contractual text and coordinate long-running workflows; (ii) a document-AI and knowledge-graph stack, which transforms unstructured contracts and field evidence into machine-queryable state; and (iii) a cloud infrastructure and data pipeline layer, which ensures the previous two layers can operate at portfolio scale under the security and compliance constraints of the energy sector. This chapter describes each layer at the level of detail required to assess the technical feasibility of the business plan presented in subsequent chapters.

2.1 Agentic AI foundations

Agentic AI refers to a class of systems in which a large language model (LLM) does not merely produce a single text response but acts as a planner that decomposes a goal into sub-tasks, selects tools to execute each sub-task, observes the tool outputs, and iterates until the goal is satisfied [27, 28]. The three architectural primitives that underpin modern agentic systems are function calling, structured memory, and multi-agent orchestration.

Function calling — introduced as a first-class capability by frontier LLM providers in 2023 — allows the model to emit a JSON-formatted intention to invoke an external tool (a database query, an API call, a deterministic parser) instead of generating free text. This primitive is the bridge between probabilistic language reasoning and deterministic software systems, and it is what makes agents auditable: every action the model takes leaves a structured trace that can be logged, replayed, and reviewed. The more recent Model Context Protocol (MCP) standardises the way tools and data sources are exposed to models, which reduces integration effort when a platform must expose dozens of capabilities — the situation ConCoord is in, where a single workflow may touch contract storage, operational telemetry, ticketing, e-signature, and email.

Memory in an agentic system is deliberately split into two tiers. Short-term memory is the conversation buffer and tool-response scratchpad that the model consumes within a single run. Long-term memory is the external store — typically a relational database, a vector index, or a knowledge graph — into which the agent writes durable facts. For a contract-management platform this separation is critical: the reasoning itself is ephemeral, but the extracted obligations, deadlines, and evidence must persist deterministically and survive model upgrades.

Orchestration patterns define how multiple agents cooperate. Three patterns are relevant for ConCoord. Sequential orchestration chains specialised agents (for example: extractor → classifier → validator) in a fixed order and is appropriate when the workflow is well understood. Parallel orchestration fans out independent sub-tasks — for example, extracting obligations from the main agreement, its annexes, and all referenced standards simultaneously — and is the dominant pattern at Blueprint design time. Hierarchical orchestration places a supervisor agent on top of specialised workers, which matches the runtime monitoring loop where a coordinator agent receives asset-state events and dispatches the appropriate compliance worker.

The decision to build ConCoord on an agentic paradigm rather than on a traditional rules engine or a single monolithic LLM call is driven by three properties of contractual work. First, contracts are heterogeneous — no two FIDIC Silver Books [20] are executed with identical Particular Conditions, so the platform must reason about novel clauses rather than match templates. Second, obligations are cross-referenced — a deadline in a supply agreement may depend on a milestone defined in an interface agreement with a third party — which is naturally expressed as a multi-agent reasoning problem. Third, enforcement involves long-running workflows with human-in-the-loop checkpoints (legal review, counterparty response, evidence upload), and agentic frameworks are designed for precisely this pattern of partial automation with human oversight.

2.2 Knowledge graphs and the document-AI stack

The output of the Blueprint pipeline is not a pile of extracted text — it is a knowledge graph whose nodes are parties, obligations, deadlines, documents, assets, and events, and whose edges are the typed relationships defined by the underlying contract ontology (triggers, depends-on, references, satisfies). This design follows the broader shift in legal NLP away from flat span-extraction toward structured representations that support downstream reasoning [29, 30].

Three components of the document-AI stack deserve explicit description. The first is layout-aware parsing. Energy contracts arrive as PDFs produced by very different word processors across decades, often with scanned signature pages, landscape-oriented annexes, and embedded tables of technical specifications. A pipeline that treats the PDF as pure text loses the structural cues that distinguish a clause heading from a footnote, or a defined term from a cross-reference. ConCoord combines an OCR stage with a layout-aware parser that preserves page coordinates for every token, which in turn feeds the character-level citation that underlies the platform's auditability guarantee.

The second component is retrieval-augmented generation (RAG) [31], used in two distinct ways. At Blueprint time, RAG surfaces analogous clauses from previously analysed contracts so the extraction agent can calibrate its interpretation against precedent. At runtime, RAG lets a user — or an agent acting on the user's behalf — ask natural-language questions about a portfolio of live obligations ("which O&M contracts impose availability-guarantee penalties above 5% of annual fee?") and receive answers grounded in the underlying knowledge graph.

The third component is the ontology itself. Rather than hard-coding a schema, ConCoord maintains a domain ontology — drawn from FIDIC conditions [20, 21], ICC and CMAA model contracts [22, 23], and the Spanish regulatory corpus [18, 19] — that enumerates the obligation types, evidence categories, and trigger events typical of renewable-asset contracts. The agents use this ontology as a labelling scaffold: a contractual clause is first classified against the ontology, then, if no category fits, a new node is proposed and flagged for human curation. This hybrid approach gives the system a stable vocabulary without forcing every contract to be squeezed into a rigid template.

2.3 Cloud infrastructure and data pipeline

ConCoord is deployed as a multi-tenant SaaS on Amazon Web Services, with the architectural goal of isolating customer data at the storage layer while sharing the compute layer for cost efficiency. The core runtime runs on Amazon ECS with Fargate, which eliminates the need to manage EC2 capacity and aligns compute cost with request volume — an important property during the first-year ramp, when monthly active contracts are in the tens rather than the thousands.

The persistent layer relies on three services. Amazon RDS for PostgreSQL stores the knowledge graph and all operational metadata, with per-tenant schemas and row-level security to enforce isolation. Amazon S3 stores the raw contract PDFs and uploaded evidence documents, with server-side encryption using customer-managed KMS keys so that even a full AWS-account compromise would not expose cleartext documents. Amazon OpenSearch

— operated as a managed vector index — holds the embeddings used by the RAG pipeline. Each of these services is deployed in the eu-west-1 (Ireland) region to keep data inside the European Union by default, with an eu-south-2 (Madrid) failover option on the roadmap for customers whose counterparty contracts require Spanish-soil residency.

The LLM layer is intentionally vendor-agnostic. A thin internal gateway speaks to Amazon Bedrock for hosted Anthropic and Meta models and to the OpenAI API for frontier models; the gateway enforces rate limits, logs every prompt and completion for audit, and strips personally identifiable information before forwarding requests. This abstraction means ConCoord can route workloads to whichever provider offers the best quality-price point at a given moment — a non-trivial economic lever given that inference prices for frontier-class models have fallen by more than 80% per token since 2024.

Asynchronous orchestration is handled by Amazon SQS and Step Functions. Long-running Blueprint builds, nightly portfolio re-scoring, and scheduled obligation checks are all modelled as state machines rather than as direct API calls, so that transient LLM outages or tool failures trigger automatic retries with exponential backoff instead of propagating to the user interface. Identity is federated through Amazon Cognito with enterprise SAML so that customer single-sign-on integrates without bespoke engineering, and egress is fronted by Amazon CloudFront with a Web Application Firewall configured against the OWASP top-ten ruleset.

From a compliance posture, this infrastructure supports the controls required by ISO 27001 and SOC 2 Type II — the two certifications that procurement teams at Spanish IPPs typically request during vendor onboarding — and positions the platform for the more demanding information-security requirements that apply to critical-infrastructure operators under the NIS2 Directive transposition.

Chapter 3. STATE OF THE ART

3.1 Existing commercial solutions

Contract-lifecycle-management (CLM) software has matured into a distinct SaaS category over the past decade, growing from under $1B in 2020 to a projected $5.4B by 2036 [26]. Five vendor archetypes dominate the landscape, and understanding the gaps between them is essential to positioning ConCoord.

Enterprise CLM suites — Icertis, SirionLabs, and Conga CLM — target multinational corporations with very large counterparty bases (thousands of NDAs, MSAs, and SOWs). Their strengths are workflow governance, e-signature integration, and repository management. Their weakness, from the perspective of an infrastructure operator, is that they treat contracts as documents to be tracked rather than as live obligations to be enforced against an external physical reality. None of the three exposes a primitive for "asset state" or connects contractual clauses to operational telemetry.

Mid-market CLM platforms — DocuSign CLM, Ironclad, LinkSquares, and Evisort — emphasise ease of deployment and reasonable out-of-the-box AI extraction. Ironclad and Evisort, in particular, invested early in LLM-based clause classification and now ship pre-trained extractors for common contract types. These platforms are a natural fit for sales and procurement workflows but, like the enterprise suites, stop at the edge of the document — they do not model the asset-level consequences of what they extract.

Point solutions have emerged around specific bottlenecks in the construction and infrastructure segment. Clearstory [24, 25] focuses on change-order management between general contractors and subcontractors; Summize [33] targets the in-house legal review workflow with an LLM-powered negotiation copilot. These tools are the closest analogues to ConCoord in terms of AI-native approach, but each addresses a narrow slice of the contract lifecycle — change orders or legal review — rather than the full span from extraction to enforcement.

The fourth archetype is BIM and project-management software — Procore, Autodesk Construction Cloud, Oracle Primavera — which owns the operational and schedule data against which contractual obligations should be evaluated. These platforms contain the raw material of enforcement (progress, delays, defects) but are blind to the contract text, and their vendors have historically outsourced the legal layer to partners.

The fifth archetype, finally, is the consultancy-plus-spreadsheet combination that dominates the actual day-to-day reality of Spanish renewable IPPs: contract-and-risk consultancies reviewing executed contracts on an ad-hoc basis, feeding their findings back into asset-owner Excel registers maintained by asset managers. This is the pattern that ConCoord most directly replaces.

The comparison below summarises how the archetypes distribute across the four capabilities that define the ConCoord wedge.

Vendor archetype	Document repository	AI extraction	Asset-state linkage	Enforcement workflow
Enterprise CLM (Icertis, Sirion, Conga)	Deep	Moderate	Absent	Generic approvals
Mid-market CLM (DocuSign, Ironclad, Evisort)	Deep	Strong	Absent	Generic approvals
Point solutions (Clearstory, Summize)	Narrow	Strong within scope	Narrow within scope	Narrow within scope
BIM / project management (Procore, Primavera)	Absent	Absent	Deep	Operational only
Consultancies + Excel	Shallow	None	Manual	Reactive
ConCoord	Deep	Strong	Deep	Contract-native, automated

No incumbent currently combines contract-text depth with asset-state awareness inside a single workflow. This is both the market gap and the source of ConCoord's defensibility.

3.2 Relevant academic research

Academic work on automated contract analysis has accelerated in step with the broader rise of large language models. Four strands of research frame the ConCoord design decisions.

The first strand is legal-NLP benchmarking. The CUAD dataset of commercial contracts [30] and the LegalBench evaluation suite have established that clause-type classification and span extraction on contracts are tractable tasks for modern LLMs, with top systems now exceeding 85% F1 on well-defined clause categories. The open question, which directly motivates the Blueprint architecture described in Chapter 5, is how to move from clause classification to obligation-level semantic modelling — capturing not only that a clause belongs to a given category, but what behaviour it requires, from whom, and under what conditions.

The second strand is retrieval-augmented generation [31]. RAG addresses the well-documented tendency of LLMs to hallucinate when asked questions that require specific factual grounding. In the contract domain, RAG is doubly useful: it grounds answers in the exact clause text (essential for auditability), and it lets the platform scale to portfolios that far exceed any model's context window. ConCoord's runtime Q&A layer is a direct application of this line of research.

The third strand is agentic benchmarks and safety evaluations [27, 28, 29]. Recent work on function-calling reliability, multi-step reasoning stability, and refusal-and-escalation behaviour provides the empirical backbone for the claim — made repeatedly in this TFG — that agentic systems can be deployed in high-stakes workflows if they are constrained by deterministic tool interfaces, human checkpoints, and complete audit logs. The ConCoord orchestration design explicitly inherits the supervisor-with-specialised-workers pattern validated in this literature.

The fourth strand is the intersection of smart contracts and traditional contract law [32]. Although ConCoord is not a blockchain product, the conceptual lineage from Szabo's early formulation of contractual terms as executable code remains relevant: the platform implements the programmable-obligation vision without requiring parties to move execution on-chain, which matches the practical reality that Spanish renewable IPPs and their counterparties contract under Spanish civil law, not under smart-contract jurisdictions.

Taken together, this body of research suggests that the technical feasibility of an agentic contract-management platform is no longer in question; the open questions are in systems design, domain modelling, and go-to-market execution, which is precisely where this TFG makes its contribution.

3.3 Limitations of current approaches

The management of contractual obligations in infrastructure projects is today handled by a fragmented ecosystem grouped into three categories: contract-and-risk-management consultancies, specialized law boutiques, and contract-lifecycle-management (CLM) or BIM software platforms. Analysis of 112 relevant companies across these segments reveals an industry that remains overwhelmingly manual, relationship-driven, and reactive — structurally unable to keep pace with the complexity and scale of modern renewable-energy portfolios.

Consultancies are typically engaged after a problem has materialised: a delay has occurred, a claim has been filed, or a dispute has escalated. Their value proposition is forensic — reconstructing what went wrong and quantifying damages after the fact. Proactive, real-time monitoring of contractual compliance across an entire portfolio is economically unfeasible under a billable-hours model.

Law boutiques specialise in infrastructure, energy, and construction law. Their involvement typically begins at the escalation point, when a contractual issue has become a formal claim or dispute. While invaluable for high-stakes disputes [21, 22], this model does nothing to prevent the contractual misalignment that triggers those disputes in the first place.

CLM and BIM platforms digitise document storage and workflow approvals but do not bridge the gap between the physical state of an infrastructure asset and its contractual framework. Clearstory [24, 25] is the most relevant digital competitor — narrow in scope (change-order management). Traditional CLM vendors such as Icertis, Agiloft, and ContractPodAi store a FIDIC Silver Book [20] but cannot interpret whether a delay event triggers a time extension under Sub-Clause 8.5. The global CLM market is forecast to grow from $1.8B in 2026 to $5.4B by 2036 [26], confirming the category trajectory, but none of the incumbents closes the physical-state/contract gap.

3.4 Differential contribution of the project

ConCoord transforms contracts from passive documents into active compliance assets. Four capabilities define its differential contribution:

Deadline prevention. Every time-bound obligation is pre-scheduled at instantiation, with graduated reminders that escalate as deadlines approach — no obligation can be silently missed.

Evidence-grounded compliance. Every extracted element carries a character-level citation back to the source clause, and every fulfilled obligation is backed by an uploaded document linked to its tracker.

Idiosyncrasy-tolerant modelling. The agents treat the underlying ontology as a labelling scaffold, not a rigid template, so non-standard drafting and compound obligations are represented faithfully.

Full auditability. Every state transition, document upload, notification, and agent decision is appended to a chronological event log, producing a complete timeline from clause text to enforcement action.

Together, these capabilities are the ConCoord wedge: bridging the physical-state ↔ contract gap in real time, which no incumbent in the three current categories delivers.

Chapter 4. DEFINITION OF THE WORK

4.1 Justification

The rationale for this TFG rests on three reinforcing arguments: a technical argument about why agentic AI is the right tool for this problem at this moment, a research argument about the academic gap that the work addresses, and a social-industrial argument about why the Spanish renewable sector is both the right first market and a meaningful lever on the energy transition.

4.1.1 Technical justification

Until the most recent generation of LLMs, building a system that could read an arbitrary energy contract, reason about obligations cross-referenced across annexes and third-party agreements, and orchestrate long-running enforcement workflows was not technically feasible at a cost compatible with mid-market SaaS pricing. Two developments have changed that equation. First, frontier LLM performance on legal-NLP benchmarks now exceeds the threshold at which domain-fluent professional review becomes the bottleneck, not model quality [29, 30]. Second, inference prices for frontier-class models have fallen by roughly an order of magnitude in twenty-four months, which turns what was a research-grade prototype in 2023 into a commercially viable product in 2026. Agentic orchestration [27, 28] provides the architectural framing needed to turn that raw capability into an auditable, human-supervised workflow — which is the level of reliability a legal-grade product requires.

4.1.2 Research / academic justification

The TFG sits at the intersection of three research communities that rarely publish together: legal NLP [30], agentic AI and LLM tool-use [27, 28, 29], and the systems-engineering literature on infrastructure contract administration [20, 21, 22, 23]. Each community has produced relevant pieces — clause classifiers, multi-agent frameworks, obligation taxonomies — but the integration of the three into a production-oriented architecture targeting a specific industry vertical is an under-explored design space. The contribution of this TFG is to propose that integration explicitly, validate it against the requirements of Spanish renewable IPPs, and show through the pricing and unit-economics analysis of Chapters 6 and 7 that it is commercially sustainable.

4.1.3 Social and industrial justification

Spain is one of the most attractive renewable-generation markets in Europe, with over 36 GW of installed solar PV capacity and an additional 15+ GW in advanced development [1, 15]. The pace of deployment required to meet national and EU 2030 targets [18] means that operational efficiency — not just capacity construction — has become the binding constraint. A meaningful fraction of the value that those projects are meant to deliver is lost inside the contractual layer: delayed claims, missed penalty windows, and opaque O&M performance obligations that prevent financiers from pricing risk accurately [3, 4, 5]. Any technology that recovers part of that lost value translates directly into a lower cost of capital for new renewable build-out, which is the point of contact between a business-plan TFG and the energy-transition challenge.

Beyond the energy angle, the industrial justification is that mid-market Spanish IPPs are structurally underserved by enterprise CLM vendors and cannot afford the consultancy-led approach that would in principle provide equivalent coverage. This creates a wedge for a product that is purpose-built for the segment and priced to match its unit economics.

4.2 Objectives

The TFG has a single general objective and a set of specific objectives that together constitute the body of the work.

4.2.1 General objective

To design, justify, and evaluate the commercial viability of ConCoord, an agentic software platform that continuously links the physical state of infrastructure assets to their full contractual framework, with Spanish renewable Independent Power Producers as the beachhead market.

4.2.2 Specific objectives

The general objective is decomposed into seven specific objectives:

1. Characterise the contractual-administration pain of Spanish renewable IPPs and size the addressable market in terms of active plants, contracts under management, and yearly enforcement events.

2. Review the state of the art — both commercial (CLM, BIM, point solutions) and academic (legal NLP, agentic AI, RAG) — and articulate the differential contribution of ConCoord against it.

3. Specify the solution architecture at a level of detail sufficient to estimate build cost, infrastructure cost, and runtime unit economics.

4. Design three alternative pricing strategies, compare them against a common unit-economic model, and recommend a go-to-market sequence consistent with the commercial realities of the target segment.

5. Project a five-year P&L under the recommended strategy, including ramp assumptions, gross-margin evolution, and steady-state EBITDA margin.

6. Produce a risk register covering product, commercial, legal, and execution risks, with mitigation approaches aligned to the ICAI-ICADE risk-management framework.

7. Align the project to the UN Sustainable Development Goals — principally SDG 7, 9, 12 and 13 — and articulate the contribution pathway.

4.3 Methodology

The research and design process followed a sequential-but-iterative methodology structured around five steps.

The first step was a literature review covering agentic AI, retrieval-augmented generation, legal NLP, Spanish and EU energy regulation, and the major model-contract frameworks (FIDIC, ICC, CMAA). The purpose of this step was to build a vocabulary sufficient to communicate across the technical, legal, and industrial audiences of the thesis.

The second step was a competitive and stakeholder analysis. Public information on 112 potentially relevant vendors was gathered and classified into five archetypes, and informal interviews were conducted with asset managers at mid-size Spanish IPPs to validate the pain points observed in the secondary research.

The third step was the architectural specification. The Blueprint / Runtime split described in Chapter 5 was drafted, iteratively challenged against the cost model, and revised when the initial assumptions — for example on per-contract LLM spend at Blueprint time — proved incompatible with the pricing model targeted at mid-market IPPs.

The fourth step was the pricing and financial modelling. Three alternative pricing strategies were built against a shared unit-economic model, and the five-year P&L was stress-tested under plausible ramp-rate scenarios. The recommendation in Section 6.4 emerged from the comparison rather than being assumed upfront.

The fifth step was risk identification and SDG alignment, carried out in parallel with the financial modelling to ensure that risks with quantitative impact (for example LLM cost inflation or regulatory change) were reflected in the sensitivity analysis and not treated as narrative afterthought.

The process was iterative in the sense that each review session with the thesis director led to revisions that propagated across multiple earlier steps — for example, a refinement of the target customer profile in the commercial step triggered updates to both the market-sizing section and the pricing-strategy comparison.

4.4 *Temporal planning*

The Gantt chart in Figure 1 is the single visual summary of the first commercial quarter. Each of the four workstreams — Product & Engineering, Commercial, Legal-Compliance-Corporate, and Operations & Team — runs continuously across all three months; the monthly boundaries mark shifts in emphasis rather than handoffs. Three anchor milestones define the quarter: signed LOI in Week 4, first pilot live in Week 5, and first paid annual contract by Week 12.

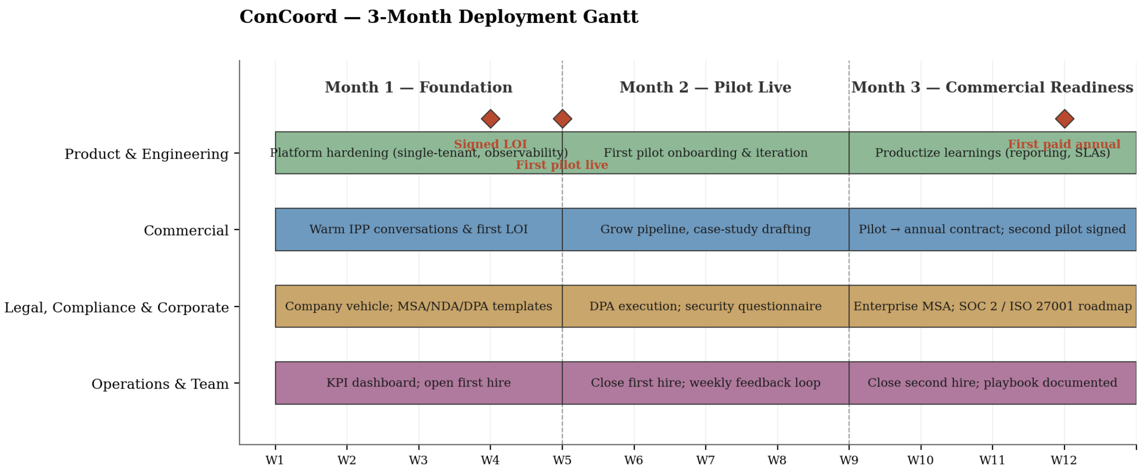

Figure 1 — ConCoord 90-day deployment Gantt: four parallel workstreams and three anchor milestones.

4.5 *Commercialisation and scalability potential*

ConCoord's commercial trajectory is designed along two complementary axes: a geographic axis (from Spain into other European renewable markets such as Portugal, Italy, Germany, and Poland, where many Spanish IPPs already operate) and a sectoral axis (from renewable energy into data centres, conventional energy, transportation, and water/utilities, all of which share the same contractual architecture). In the base case, the platform reaches 230 active plants in Year 5, generating $11.0M in Annual Recurring Revenue with EBITDA margin above 80%. Section 7 details the underlying unit economics and P&L assumptions.

Scalability is supported by the separation of design-time and runtime pipelines: once a Blueprint is produced for a given contract template, the runtime cost of monitoring additional signed instances is almost entirely variable. This property — combined with falling frontier-LLM inference prices — drives gross margin expansion from 75% in Year 1 to 91% by Year 5.

Chapter 5. PRODUCT AND SOLUTION ARCHITECTURE

5.1 High-level architecture

At the architectural level the system is organised around a single pivotal artefact — the Blueprint — that is produced by one pipeline and consumed by another. The design-time pipeline performs a one-time, compute-intensive analysis of a contract and its clauses to produce a structured, annotated state machine. The runtime pipeline treats that Blueprint as an immutable template and spawns independent, long-running monitoring processes from it. This clean separation is what allows the analytical and the operational concerns of the system to evolve independently: the modelling agents can be upgraded without disturbing live instances, and the runtime semantics can be refined without re-analysing signed contracts.

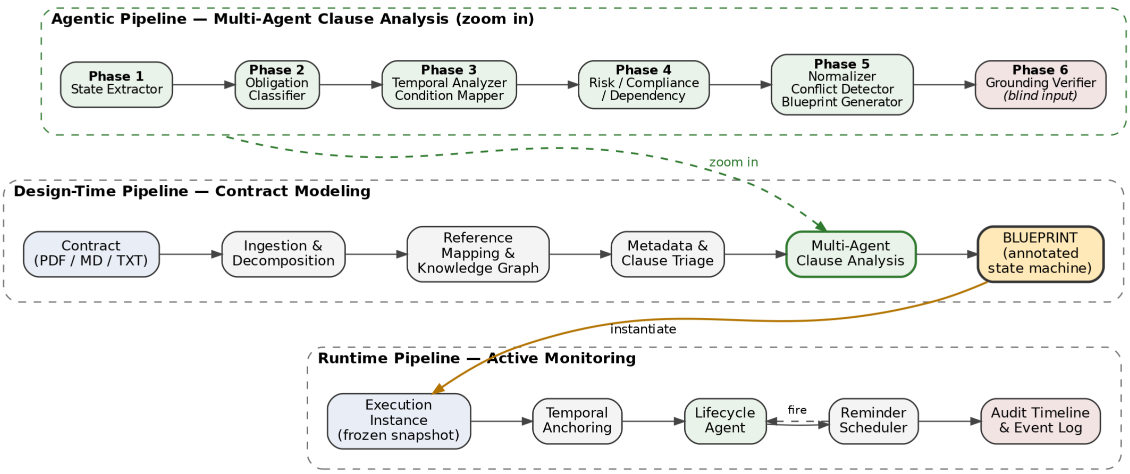

Figure 2 — End-to-end architecture. The Blueprint is the immutable hand-off between the design-time and runtime pipelines.

Key insight: the Blueprint is both the final output of the analysis pipeline and the sole input to the monitoring pipeline. Everything upstream of it is about understanding contracts; everything downstream is about executing them. This architectural pivot is what lets the system combine deep LLM-based semantic analysis with deterministic, auditable runtime behaviour without the two concerns contaminating each other.

5.2 Design-time pipeline: from contract to Blueprint

The design-time pipeline is implemented as a configurable step chain in which each stage declares its inputs, outputs, and dependencies. Steps support caching — a clause that has already been analysed is not re-analysed unless its inputs change — and lazy initialisation, which makes the pipeline well suited for batch processing large corpora. The following subsections describe each stage in order.

5.2.1 Ingestion and hierarchical decomposition

The system accepts contracts as PDF, Markdown, or plain text. PDFs — whether digital or scanned — are routed through a neural document parser (LlamaParse) that converts them into structured Markdown while preserving heading hierarchy, numbered sections, and prose flow. Text and Markdown files skip the conversion stage and enter the decomposition logic directly.

Once a document is in Markdown form, a decomposition step segments it into discrete clause units. Two complementary strategies are applied in priority order: first, header-based splitting at Markdown heading boundaries, which preserves the logical hierarchy of parts and sub-parts; and second, numbered-section detection as a fallback when explicit heading markup is absent. Each resulting clause carries its text content, its positional index in the document, the extracted clause number, and the heading title, along with structural metadata for traceability.

5.2.2 Reference mapping and knowledge graph

Legal contracts are richly cross-referential: clauses cite other clauses, annexes, articles, and external instruments. An LLM-based extraction step operates over each clause independently to identify these references, recording the exact source text of each reference, its type, a normalised target identifier, and a contextual description. Extraction is constrained by a JSON schema and runs bilingually in Spanish and English.

A resolution step then builds a directed knowledge graph over the entire document: each clause becomes a node, each reference becomes a directed edge, and unresolved targets (external instruments, missing annexes) are preserved as typed placeholder nodes. The resulting graph exposes structural connectivity (which clauses depend on which others), centrality signals (the most-referenced clauses, which usually correspond to the contract's core operative provisions), and isolation signals (clauses with no incoming or outgoing references, which are typically boilerplate).

5.2.3 Document metadata and clause-level triage

Before clause-level analysis begins, an LLM agent analyses a representative sample of the document's opening clauses to produce a document-level summary: contract type, parties and their roles, governing jurisdiction, key dates, and primary currency. This metadata is persisted alongside the document record and made available as shared context for every downstream agent.

A second agent performs clause-level triage. Not every clause warrants deep modelling: purely definitional, recital, or administrative clauses carry no time-bound obligations or state-changing logic, and analysing them with the full pipeline would add cost and noise without analytical value. The triage agent classifies each clause on a binary axis — does this clause require active modelling? — and when its determination is high-confidence negative, the full analysis pipeline is bypassed for that clause. This short-circuit ensures analytical depth is spent only where it adds value.

5.2.4 Multi-agent clause analysis

The core of the design-time pipeline is a sequence of twelve specialised agents organised into six phases [29, 30]. Each agent receives the outputs of its predecessors as structured context, adds a new semantic layer, and passes the enriched result downstream. All agents run at temperature 0 and produce JSON-schema-validated output. Every extracted element carries character-offset source evidence, ensuring full traceability from Blueprint back to clause text [31].

5.2.5 The Blueprint

The output of the design-time pipeline is a Blueprint: an executable, annotated state machine that models the complete lifecycle of a clause [32]. It is composed of three structural layers. States represent the contractual conditions in which the parties exist at a given moment, each typed by its lifecycle role and carrying the obligations, actors, and deadlines that apply while the contract occupies it. Transitions are directed edges representing the events or conditions that move the contract between states, typed by trigger mechanism and carrying guards and activated consequences. Consequence paths are ordered sequences of states from a point of divergence to a terminal state, giving the Blueprint a narrative structure that mirrors how legal outcomes are actually reasoned about.

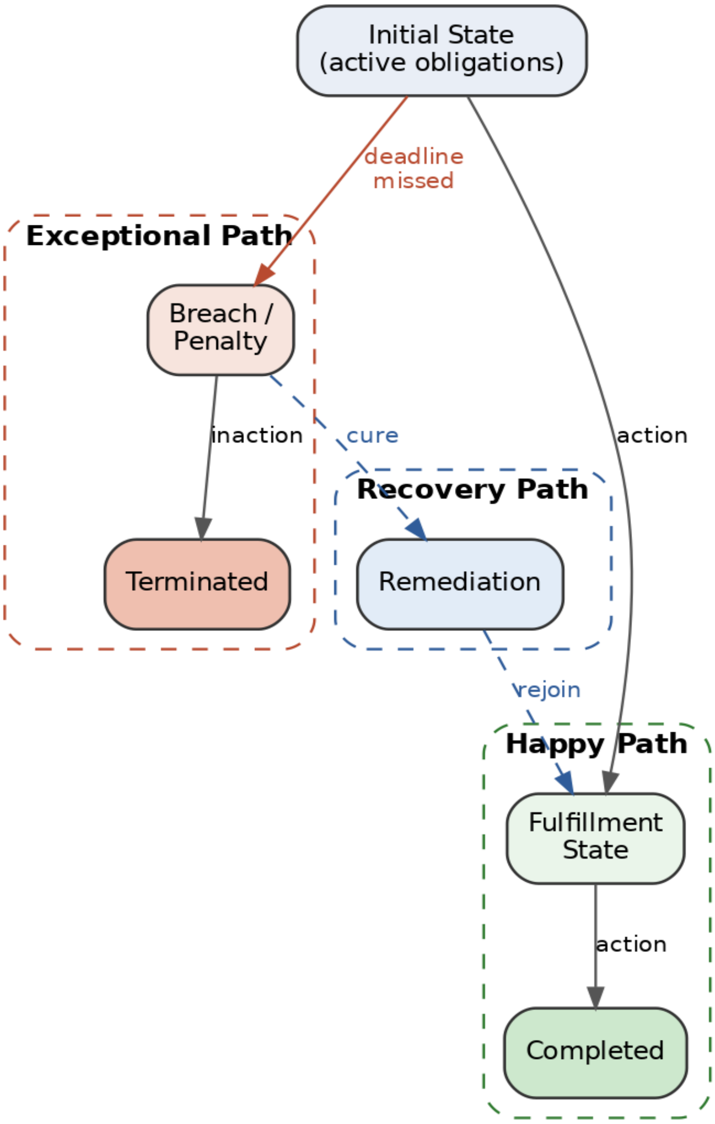

Figure 3 — Anatomy of a Blueprint. States are grouped into named consequence paths; transitions carry their trigger type and any guards.

5.3 Runtime pipeline: active lifecycle monitoring

The runtime pipeline takes a Blueprint and drives it through the lifecycle of a real, signed contract. Where the design-time pipeline is agent-heavy and analytical, the runtime pipeline is deterministic-first: the AI is deployed only where judgment is required, and the operationally critical mechanisms (deadline scheduling, state transitions, event logging) run as predictable code.

5.3.1 Execution instances and frozen snapshots

Monitoring begins with the creation of an Execution Instance. A Blueprint serves as the template; when a user triggers instantiation, the system binds it to a live runtime context containing the contracting parties, a reference date, and any clause-relevant variables. At creation time the system takes a frozen snapshot of the Blueprint and attaches it permanently to the instance. This immutability is deliberate: the Blueprint may later be refined or corrected, but a live monitoring process must reflect the state of the agreement at the moment of execution. Each instance is fully independent — multiple clauses from the same contract, or multiple contracts sharing a template, each produce an isolated monitoring process with no shared mutable state.

5.3.2 Automatic activation at instantiation

Instantiation is not a passive event. Three automatic actions are triggered without user intervention. First, temporal anchoring resolves the reference date and converts every relative temporal expression in the Blueprint ("within 30 days", "before the end of the fiscal quarter") into concrete calendar deadlines. Second, deadline pre-scheduling inspects every outgoing edge from the initial node and computes a graduated set of reminder and deadline events for each time-bound edge — computed deterministically at startup, not on demand, so no deadline can be silently missed. Third, the lifecycle agent is invoked: it reads the full Blueprint and the instance context, activates the obligations of the initial state, notifies the relevant parties, and creates pending evidence upload requests. By the time activation completes, the instance is fully live.

5.3.3 State-scoped obligations

Each node in the Blueprint carries a set of enriched obligation records describing who must do what, of what type, and by when. Obligations are state-scoped: they are activated when the instance enters the node to which they belong, not loaded globally at instantiation. At any moment, the active obligation set reflects only the duties relevant to the current phase, preventing confusion between past, present, and future obligations. Each obligation has its own tracker that progresses through a defined lifecycle — pending → in-progress → fulfilled — with overdue and waived as terminal exceptions. Fulfillment requires evidence: the lifecycle agent issues document upload requests appropriate to the obligation type (payment receipts, delivery notes, insurance certificates), and submitted documents are linked to the tracker, producing an evidentiary chain from clause text to documented compliance.

5.3.4 State machine navigation

At any point during the contract lifecycle the system maintains a precise picture of where the contract stands and where it can legally go next. From the current node, the Blueprint exposes a set of outgoing edges with their triggers, trigger types, responsible actors, guards, and legal consequences. The lifecycle agent continuously evaluates which transitions are actionable: it inspects the active obligation set, checks whether guards are satisfied (e.g., whether a required document has been submitted), and when all guards for an edge are met, executes the transition — advancing the current node, writing a transition record to the audit log, and re-activating for the new state. For dispute resolution or administrative intervention, an editor mode allows authorised users to force a transition bypassing guard evaluation; such overrides are explicitly logged as manual interventions.

5.3.5 The reminder scheduler

Time management is handled by a purpose-built, deterministic Reminder Scheduler that operates independently of the AI agent. Every time the system enters a node, the scheduler inspects each time-bound outgoing edge and computes its deadline. For each deadline it creates a graduated set of reminder events: the closer the deadline, the more frequent the reminders. A deadline two years away generates reminders at 730, 365, 180, 90, 30, and 7 days; a deadline ten days away generates reminders at 7, 3, and 1 day. Edges whose expiry has direct contractual consequences receive additional last-minute reminders regardless of total duration. When a node is exited, all pending reminders for the departed node are cancelled and the entering node's reminders are freshly scheduled, ensuring the active reminder set always reflects the current contractual position.

A background scheduler polls the event store at short intervals and fires events as they fall due. Depending on the event's action type, firing may automatically execute an edge transition (for inaction deadlines, where the passage of time itself is the trigger), send an escalation notification, or invoke the lifecycle agent for re-planning. Every event — activation, fulfilment, document request, notification, timer firing, agent decision — is appended to a chronological event log that constitutes the complete audit timeline.

5.3.6 Consequence-aware enforcement

Key insight: because each transition in the Blueprint carries explicit consequence annotations, and because consequence paths classify outcomes by their legal nature, the system can reason about the gravity of an approaching deadline, not just its proximity. When an obligation approaches its deadline without fulfilment, the graduated reminder system escalates urgency; when the deadline passes, the inaction edge fires automatically, transitioning the instance into the corresponding consequence state — a penalty state, a dispute state, or a termination state — and the lifecycle agent activates the obligations that state carries. This closes the loop from clause text to live enforcement action.

5.4 Cross-cutting architectural principles

Ontology as scaffold, not template. A central design stance is that contracts are idiosyncratic: even within the same type and jurisdiction, no two clauses share the same obligation structure or consequence logic. Conventional template matching — "if the clause looks like X, label it Y" — fails here.

Evidence grounding end to end. Every element of the Blueprint — every state, transition, and path — is annotated with character offsets into the source clause text. Every obligation fulfilled at runtime is linked to an uploaded evidentiary document. Every state transition is recorded in the audit log with its trigger and supporting evidence. The traceability chain from clause text to enforcement action is continuous and queryable.

Deterministic where it matters, agentic where it helps. The system deliberately partitions responsibility. Semantic understanding is performed by LLM agents, because that is the problem class where they excel. Operationally critical mechanisms — deadline computation, reminder scheduling, state transitions, audit logging — are implemented as deterministic code, because their correctness must be auditable and reproducible. The two are composed, not confused.

Instance isolation and dual verification. Each Execution Instance carries a frozen Blueprint snapshot and its own state, timers, and audit trail, with no shared mutable state between instances — a correction to a template never retroactively alters a live contract. Grounding, likewise, is verified at two complementary levels: a deterministic string-matching engine checks that every piece of cited source evidence actually appears in the clause text (catching hallucinations), while the Grounding Verifier agent, operating on clause text and Blueprint alone, catches omissions that string matching cannot. The two mechanisms have complementary blind spots and together produce a robust completeness guarantee.

5.5 Technology and implementation overview

The implementation combines a neural document-parsing layer (LlamaParse) for PDF ingestion, a configurable step-chain orchestrator with per-step caching, and an agent framework that enforces JSON-schema-validated output at temperature 0. Agents operate bilingually (Spanish and English). The runtime layer is built around a persistent event store, a background scheduler, and a lifecycle agent with an expanded token budget for assembly-critical decisions. The Blueprint is rendered as a swim-lane diagram (one lane per consequence path), and persistence is append-only at the event level — enabling the audit timeline to be reconstructed at any moment.

Chapter 6. BUSINESS MODEL AND PRICING

6.1 ROI-first pricing philosophy

ConCoord's pricing is built around a single principle: every euro a client pays must be visibly and measurably offset by value delivered. This is not a software cost — it is an investment in loss prevention, claim enforcement, and contract intelligence. The value story is therefore the pricing strategy.

The infrastructure sector currently spends between €300,000 and €600,000 per year per mid-sized IPP on reactive contract management across consultancy firms, legal boutiques, and software platforms [23] and still lacks real-time visibility into its contractual exposure. ConCoord replaces this fragmented, reactive stack with a single agentic layer at a fraction of the combined cost.

The ROI case for ConCoord rests on two pillars. The first is cost substitution: replacing a significant portion of existing consultancy and legal spend with a more efficient platform. The second, and substantially larger, pillar is loss prevention: proactively identifying and enforcing contractual rights that are routinely left on the table or forfeited entirely under the current reactive model.

A single missed force-majeure notification window under a FIDIC EPC contract [20] can forfeit a €3–5 M claim. An O&M provider who knows their client is not tracking availability metrics precisely can quietly under-report by 1–2 percentage points — worth hundreds of thousands annually on a 200 MW plant. PPA counterparties routinely calculate curtailment, negative-price provisions, and P50 deviations in their own favour when there is no automated counter-check [6]. ConCoord closes each of these gaps in real time.

6.2 Client ROI: quantified value drivers

For a representative mid-sized IPP managing a 200–300 MW portfolio across 6–8 plants, ConCoord protects value across nine measurable dimensions. The estimates below have been recalibrated against the industry data sources cited in Chapter 1:

Value Driver	Annual Exposure	Prev. Rate	Value Protected
Consultancy cost replacement [23]	$300K–$600K/yr	70%	$200K–$420K
Legal dispute costs avoided [5]	$800K/dispute	65%	$520K

EPC disputed claim value [4]	$835K exp./yr	65%	$543K
O&M SLA / availability drift [4]	$70K (3.5% of $2M O&M)	80%	$56K
PPA counterparty risk (curtailment / neg. price)	$300K (2% of $15M)	75%	$225K
Missed warranty / insurance / guarantee claims [5]	$500K (1% of $50M asset)	80%	$400K
Unenforced judgments / stale awards [5]	Client-specific	80%	Variable
EPC liquidated damages not enforced [20]	$200K (2% of $10M EPC)	70%	$140K
Missed force-majeure / time-extension rights [3]	$100K (1% of $10M EPC)	80%	$80K
TOTAL VALUE PROTECTED (representative 8-plant portfolio)			~$2.16M/yr

Against a ConCoord Strategy 2 cost of $403K/yr for an 8-plant portfolio, the net annual benefit is approximately $1.76M/yr, a ~5.4× ROI multiple.

6.3 *Pricing strategies*

ConCoord is offered under three complementary pricing strategies. Each maps to a different buyer archetype and ARR profile; together they form a flexible commercial menu.

6.3.1 Strategy 1 – Per-plant pricing

Structure: $40,000 one-time onboarding fee per plant, plus a $5,000 annual maintenance fee per active plant.

Parameter	Amount	Rationale
Onboarding fee (per plant)	$40,000	One-time: contract ingestion, data integration, obligation mapping
Annual maintenance (per plant)	$5,000	Ongoing: monitoring, alerts, updates, support
Typical 8-plant portfolio — Year 1	$360,000	$40K × 8 onboarding + $40K maintenance
Typical 8-plant portfolio — Year 2+	$40,000/yr	Maintenance only; near-zero cost of ownership

ROI case. For an 8-plant client, the Year 1 investment of $360,000 compares against a current annual consultancy spend of $400K–$600K [23]. From Year 2, the $40,000 maintenance cost represents a 90%+ reduction in the platform cost component of contract management, while delivering real-time coverage no consultancy firm can match. Against $2.16M in annual value protected, Year 1 ROI is 6.0× and Year 2+ ROI exceeds 54×.

Strategic note. The $5,000 annual maintenance fee is intentionally conservative for market entry; as demonstrated ROI accumulates, this should be revised upward to $12,000–$18,000/plant/year. Strategy 1 is best positioned as the entry option for budget-conscious IPPs or for clients in early-stage purchasing conversations.

6.3.2 Strategy 2 – SaaS subscription with volume tiers

Structure: monthly subscription fee per active plant with volume-based tier discounts. No upfront commitment. Fully recurring revenue model.

Tier	Plant Range	Monthly Rate / Plant	Annual Rate / Plant
Tier 1	1 – 3 plants	$5,000 / mo	$60,000 / yr
Tier 2	4 – 8 plants	$4,200 / mo	$50,400 / yr
Tier 3	9 – 15 plants	$3,500 / mo	$42,000 / yr
Tier 4	16 – 25 plants	$2,800 / mo	$33,600 / yr
Tier 5	26+ plants	$2,200 / mo	$26,400 / yr

Market context. Enterprise CLM software costs $50K–$400K+ per year and addresses general contract storage rather than infrastructure-specific obligation monitoring. ConCoord's Tier 1 pricing of $60K/plant/year sits at the low end of this range while delivering domain-specific intelligence those platforms cannot provide. The global CLM market is forecast to grow from $1.8B in 2026 to $5.4B by 2036 at 12% CAGR [26], validating the category trajectory.

ROI case. For a mid-market IPP with 8 plants at Tier 2 pricing, annual cost is $403,200 — comparable to or lower than current consultancy spend [23]. Against $2.16M in annual value protected, this represents a 5.4× ROI. The volume-discount structure creates compounding stickiness: once a client reaches Tier 3 or 4, the per-plant cost advantage makes switching economically unattractive.

Strategic note. Strategy 2 is the primary long-term model. It generates predictable ARR, commands the highest investor valuation multiple, and scales with client portfolio growth. The recommended motion is to land clients via Strategy 3 and convert them to Strategy 2 in

Year 2, at which point ConCoord has demonstrated ROI and all contracts are already ingested.

6.3.3 Strategy 3 – Consultancy-level managed service

Structure: annual engagement fee of $200,000, positioning ConCoord as a full contract-management partner at 50% of market consultancy rates, with gross margins of ~72% vs. ~45% for traditional consultancy firms.

Parameter	ConCoord	Traditional Consultancy
Annual engagement fee	$200,000	$350,000 – $600,000+
Gross margin	~72%	~45%
Monitoring frequency	Real-time / continuous	Monthly / quarterly reports
Portfolio coverage	100% of contracts	Selective / priority items
Proactive alerting	Yes — automated	No — reactive only
Scalability	High (agentic)	Low (headcount-bound)

ROI case. Energy-sector research [5] shows that the average energy company spends $41M per year on litigation, and 35% have material awards they have not collected. Strategy 3 speaks directly to CFOs and General Counsel: same service quality, better technology, half the price, with a platform that prevents the disputes those $41M/year are spent resolving. The first-year saving of $150K–$400K in consultancy costs alone covers the engagement fee before any claim-prevention value is counted.

6.4 Strategy comparison and recommendation

The three pricing strategies examined in Sections 6.1–6.3 are not mutually exclusive; they describe the principal commercial motions available to a platform of ConCoord's shape, each with a different trade-off between time-to-revenue, contract size, and procurement friction. The recommendation that emerges from the unit-economic comparison is a sequenced deployment of the three, calibrated to the evolution of the product and the maturity of the buyer relationship.

Strategy 3 — per-contract platform fees — is the recommended land motion for the first twelve months. It keeps the annual contract size small enough to sit below the procurement-authority ceiling of an asset manager or operations director (typically €50,000–€100,000), which in turn compresses the sales cycle from six-plus months to the eight–twelve weeks targeted in the 90-day deployment plan. The per-contract structure also aligns the vendor with the customer on incremental adoption: the first signed contract delivers immediate value, and additional contracts are sold as the platform proves itself inside the account. This is the go-to-market pattern that the three anchor milestones of Figure 1 are designed around.

Strategy 2 — bundled platform-plus-usage pricing — is the recommended expansion motion from Year 2 onwards. Once a customer has three or more live plants and a stable relationship with the platform, the conversation shifts from "prove this works" to "scale this across the portfolio", and the per-contract model becomes procurement overhead. A bundled subscription priced against plant count and contract count absorbs that overhead and, critically, lets the customer commit to multi-year terms that improve net revenue retention and cash efficiency. The transition from Strategy 3 to Strategy 2 is managed as a contract conversion at the next annual renewal, not as a rip-and-replace.

Strategy 1 — a pure usage-based model — is reserved for two specific situations. The first is the budget-constrained mid-market IPP that cannot justify a fixed platform fee but needs contract-administration support; here, usage pricing provides a low-commitment entry point with the option to migrate to Strategy 2 as portfolio scale grows. The second is the ad-hoc enforcement engagement — typically a delay-damages or availability-guarantee dispute flagged by the customer after the fact — where the value is so tied to the specific claim that a platform-subscription model does not fit. In both cases, usage pricing functions as a satellite offering around the core subscription business rather than as the main commercial channel.

A fourth element of the pricing architecture, layered on top of the Strategy-2 subscription from Year 2 onwards, is an outcome-share component on successfully enforced claims. The proposed structure is 5–8% of the economic value of each enforced claim, capped per incident so that a single large claim does not distort the relationship. This component is deliberately introduced only after the platform has accumulated enough execution history to make outcome measurement credible, and it is designed to align ConCoord with the share of value the customer recovers rather than with the number of tickets processed.

The net effect of this sequenced approach is that the recommendation in this TFG is not a single price list but a commercial playbook: land with Strategy 3, expand to Strategy 2, deploy Strategy 1 for specific edge cases, and layer the outcome-share component on top once the base subscription relationship is stable. The five-year P&L of Chapter 7 assumes this sequence and sizes the ramp accordingly.

Chapter 7. FINANCIAL MODEL AND UNIT ECONOMICS

7.1 Unit economics — key assumptions

Every cost in this model traces back to three atomic inputs: the cost of processing a single contract clause, the number of clauses in a typical contract, and the number of contracts per plant. These inputs drive both the onboarding cost (incurred once) and the ongoing monitoring cost (incurred annually).

Assumption	Value	Source / Rationale
LLM cost per clause	$2.00	GPT-4o-class API; expected to decline ~20%/yr as frontier models commoditise
Average clauses per contract	25	Based on standard EPC / PPA / O&M contracts (20–30 clause range)
Contracts per plant	50	EPC, warranties, O&M, PPA, grid connection, insurance, permits
Onboarding cost per plant	$2,500	$2 × 25 clauses × 50 contracts — one-time at ingestion
Ongoing LLM spend per plant / yr	$2,000	$800 re-analysis on amendments and new documents + $1,200 interactive inference (queries, alert content, portfolio summaries)

The full pricing spreadsheet (ConCoord_Pricing_Model.xlsx) is included as supplementary material in Anexo II.

7.2 Cost structure

ConCoord's cost base is composed of three primary drivers — LLM inference, AWS cloud infrastructure, and people — plus a lean set of other operating expenses. All three drivers are variable or controllable, and each benefits meaningfully from scale, but none of them is flat: LLM and AWS costs grow roughly linearly with the active plant base, while unit economics improve as fixed elements (compliance, multi-AZ baseline, onboarding tooling) amortise across a larger installed base. In the base case, the blended operating model targets ~75% gross margin and a ~20% EBITDA margin loss in Year 1, crossing into profitability in Year 2 and reaching a steady-state 91% gross margin and 80% EBITDA margin in Year 5.

7.2.1 LLM and inference costs

LLM spend is split into three mechanically different workloads. (i) Onboarding inference runs once per plant on the full contract library and is the dominant one-off cost: ~$2,500 per plant, driven by token-heavy end-to-end extraction across 50 contracts and 25 clauses each. (ii) Ongoing re-analysis is triggered whenever amendments, side letters or new documents arrive, plus a scheduled monthly sweep against updated models; budgeted at $800 per plant per year, it grows with the active installed base, not new bookings. (iii) Interactive inference covers user-driven queries, alert-message generation, periodic portfolio summaries and the reminder scheduler's natural-language outputs; budgeted at $1,200 per plant per year, it scales with the number of active users and the obligation density of each Blueprint. Together, the ongoing workloads make LLM spend grow roughly linearly with active plants, while price-per-token declines of ~20% per year compress the ratio against revenue (from ~19% in Year 1 to ~7% in Year 5). ConCoord deliberately stays on frontier APIs (OpenAI, Anthropic, Amazon Bedrock) rather than self-hosting: this avoids GPU capex, eliminates model-maintenance overhead, and keeps clients on state-of-the-art extraction quality.

LLM COST ITEM	Y1	Y2	Y3	Y4	Y5
Total Active Plants	20	50	95	155	230
New Plants Onboarded	20	30	45	60	75
Onboarding Inference ($K)	$50	$75	$113	$150	$188
Ongoing Monitoring ($K)	$16	$40	$76	$124	$184
Interactive / Query Inference ($K)	$24	$60	$114	$186	$276
TOTAL LLM COST ($K)	$90	$175	$303	$460	$648

7.2.2 AWS infrastructure costs

ConCoord's cloud architecture runs entirely on AWS, chosen for its mature enterprise compliance posture (SOC 2, GDPR, ISO 27001), native integration with AI services (Bedrock), and the team's existing expertise. The stack is containerised from day one, enabling cost-efficient horizontal scaling as plant count grows. Infrastructure spend is modelled as the sum of (i) a fixed compliance-grade baseline — multi-AZ databases, WAF + Shield, CloudWatch observability, dev/staging environments, backups and encryption — that is essentially independent of plant count, and (ii) a per-plant variable component covering compute (ECS/EKS), PostgreSQL storage on RDS, S3 document retention, event processing (Lambda, SQS, Step Functions), and the Bedrock/OpenAI proxy. As the installed base grows from 20 plants in Year 1 to 230 in Year 5, total AWS spend scales from ~$28K to ~$162K —

around 5.8% of revenue at launch, compressing to ~1.8% at steady state as the fixed baseline amortises across more plants.

AWS SERVICE	Y1	Y2	Y3	Y4	Y5
Amazon ECS / EKS (containers)	$8.0	$16.0	$28.0	$42.0	$58.0
Amazon RDS (PostgreSQL)	$6.0	$11.0	$18.0	$26.0	$34.0
Amazon S3 (object storage)	$2.0	$4.5	$8.0	$12.0	$16.0
AWS Lambda (event triggers)	$1.0	$2.0	$3.5	$5.0	$7.0
Amazon Bedrock / OpenAI proxy	$2.0	$3.5	$5.5	$7.5	$9.5
CloudFront, Cognito, API Gateway	$1.5	$2.5	$4.0	$5.5	$7.0
SES/SNS, CloudWatch, WAF + Shield	$4.0	$6.0	$9.0	$13.0	$17.0
Dev/Staging + Data Transfer + Misc	$3.0	$5.0	$7.5	$10.0	$13.0
TOTAL AWS ($K)	$27.5	$50.5	$83.5	$121.0	$161.5

7.2.3 Headcount model

ConCoord operates on a 7-person ceiling through Year 3, expanding to a fuller commercial and engineering organisation from Year 4 onward as ARR justifies it. All roles are Spain-based (Madrid), taking advantage of Europe's deep pool of renewable-energy legal and engineering talent at a meaningful discount to UK or US equivalent compensation. The two co-founders take deliberately reduced salaries in Years 1–2, preserving runway for engineering and go-to-market investment.

Spanish employer Social-Security contributions add approximately 30% on top of gross salaries and are included in all total employer-cost calculations. The EUR/USD rate of 1.10 is used throughout for P&L consolidation.

ROLE	Y1 (2026)	Y2 (2027)	Y3 (2028)	Y4 (2029)	Y5 (2030)
Co-Founder / CEO	€40K	€50K	€65K	€80K	€90K
Co-Founder / CTO	€40K	€50K	€65K	€80K	€90K
DevOps Engineer (Sr.)	€70K	€75K	€80K	€90K	€100K
Full-Stack Engineer	€58K	€65K	€72K	€80K	€88K
Legal / Contracts Specialist	€55K	€60K	€65K	€72K	€80K
Forward-Deployed Engineer	—	€55K	€62K	€70K	€78K
Robotics Engineer	—	€48K	€55K	€62K	€70K
Total Gross Salaries (€K)	€263K	€403K	€464K	€534K	€596K
Employer SS 30% (€K)	€79K	€121K	€139K	€160K	€179K
Total Employer Cost (€K)	€342K	€524K	€603K	€694K	€775K
Total Employer Cost ($K ×1.10)	$376K	$576K	$664K	$764K	$853K

7.3 Five-year P&L summary

The P&L below combines all cost drivers against the base-case revenue projection (Strategy 2 — SaaS, Scenario 2: Mid-Market), using a consistent blended rate of $4,000 per plant per month. This yields ARR of $960K in Year 1, $4.6M in Year 3, and $11.0M in Year 5 (230 plants). Recognised revenue uses the mid-year average of the ARR run-rate, the standard SaaS accounting convention.

P&L LINE ITEM	Y1	Y2	Y3	Y4	Y5
ARR (end-yr)	$960K	$2,400K	$4,560K	$7,440K	$11,040K
Recognised Revenue	$480K	$1,680K	$3,480K	$6,000K	$9,240K
LLM / Inference (COGS)	($90K)	($175K)	($303K)	($460K)	($648K)
AWS Infrastructure (COGS)	($28K)	($51K)	($83K)	($121K)	($162K)
TOTAL COGS	($118K)	($226K)	($386K)	($581K)	($810K)
GROSS PROFIT	$362K	$1,454K	$3,094K	$5,419K	$8,430K
Gross Margin %	75%	87%	89%	90%	91%
Payroll & Employer Costs	($376K)	($576K)	($664K)	($764K)	($853K)
Other Operating Expenses	($84K)	($109K)	($149K)	($186K)	($222K)
TOTAL OPERATING EXPENSES	($460K)	($685K)	($813K)	($950K)	($1,075K)
EBITDA	($98K)	$769K	$2,281K	$4,469K	$7,355K
EBITDA Margin %	-20%	46%	66%	74%	80%

7.4 Milestone summary

The table below consolidates what each workstream is expected to deliver in each month of the first commercial quarter. It is the single artefact against which weekly progress is reviewed.

Workstream	Month 1 — Foundation	Month 2 — Pilot Live	Month 3 — Commercial Readiness
Product & Engineering	Platform hardened for single-tenant pilot: auth, ingestion, observability.	First IPP onboarded: real contracts live, Execution Instances monitored.	Productised pilot learnings: reporting layer and SLA posture.
Commercial	3–5 warm IPP conversations; first pilot LOI signed.	Pilot live; pipeline at 5–7 qualified conversations; case study in draft.	Pilot converts to annual contract; second pilot signed.
Legal, Compliance & Corporate	Company vehicle in place; MSA, NDA, DPA templates drafted.	DPA executed; first security questionnaire answered.	Enterprise MSA ready; SOC 2 / ISO 27001 roadmap defined.
Operations & Team	Founder-led delivery; KPI dashboard live; first hire opened.	First key hire closed; weekly customer feedback loop running.	Second hire closed; pilot playbook documented.

7.5 Key risks and dependencies

Pilot slippage. The largest risk is a pilot conversation that stalls between LOI and signed agreement. Mitigation: carry two parallel pilot tracks so no single customer is on the critical path.

Data access delays. Even with a signed pilot, access to a customer's contract corpus can be delayed by internal legal or IT review. Mitigation: identify a single data champion at LOI signing and pre-draft the DPA so legal review is a redline, not a cold start.

Hiring timing. The quality of the first two hires disproportionately shapes the next year. Mitigation: accept founder-led delivery in Months 1–2 and do not compromise on the Month 2 hire to meet a calendar date.

Regulatory and security posture. Infrastructure customers raise the bar on data protection faster than typical SaaS buyers. Mitigation: retain a Spanish-qualified legal advisor in Month 1, not Month 3, and treat the SOC 2 / ISO 27001 roadmap as a visible Month 3 deliverable.

Chapter 8. CONCLUSIONS AND FUTURE WORK

8.1 Summary of the work

This TFG has presented a business-plan-oriented design for ConCoord, an agentic software platform that addresses a specific and measurable failure mode in the Spanish renewable-energy sector: the gap between the physical state of an operating asset and the contractual framework that governs it. Starting from the industry evidence summarised in Chapter 1 and the technical foundations laid out in Chapter 2, the work has traced a complete argument — from market pain and competitive landscape, through architectural specification and pricing design, to a five-year P&L and a risk register — that together justify the platform's commercial viability.

The core claim defended across the preceding chapters is that the contractual layer of renewable-asset portfolios is simultaneously large enough in economic terms and under-served enough in vendor terms to support a purpose-built, AI-native product, and that the current generation of agentic frameworks and frontier LLMs has crossed the reliability and cost thresholds required to make such a product commercially practical at mid-market price points. The Blueprint / Runtime architecture described in Chapter 5 is the specific design answer to that claim, and the sequenced pricing strategy synthesised in Section 6.4 is the commercial answer. Chapter 7 translates both into explicit financial projections under a conservative ramp.

8.2 Contributions

Five concrete contributions emerge from the work and are the tangible deliverables of the TFG.

1. A characterisation of the Spanish renewable-IPP contract-administration problem grounded in industry evidence, which locates ConCoord in a specific and measurable segment of the energy-transition value chain rather than in a generic "CLM" category.

2. A two-stage technical architecture (Blueprint at design time, Runtime at operation time) specified at the level of detail required to cost, deploy, and audit, including the agent-orchestration patterns and the cloud-infrastructure choices that underpin them.

3. Three alternative pricing strategies built on a shared unit-economic model, a comparative assessment across the dimensions that matter for the target segment, and a recommended sequence (land with per-contract pricing, expand to bundled subscription, layer in outcome-share) calibrated to the real procurement patterns of mid-market IPPs.

4. A five-year P&L under the recommended strategy, with explicit ramp, gross-margin, and EBITDA assumptions, that makes the financial case transparent enough to be challenged on its inputs rather than on its conclusions.

5. A risk register and SDG-alignment analysis that together connect the commercial plan to the broader non-financial requirements a TFG of this kind must satisfy, without treating either as narrative afterthought.

8.3 Limitations

Several limitations apply to this work and deserve explicit acknowledgement.

The TFG is a plan, not a deployed product. The architectural and unit-economic claims defended here are internally consistent and grounded in public information, but they have not yet been validated by a running multi-tenant production system with paying customers. A pilot deployment at one or two Spanish IPPs — which is the first milestone of the 90-day plan in Figure 1 — is what will convert the present design into evidence.

The financial projections are built on explicit, documented assumptions rather than on observed data. Acquisition rates, conversion rates from Strategy 3 to Strategy 2, enforcement-event frequencies, and LLM-inference prices are all modelled from public benchmarks and comparable-company data; deviations from these assumptions would propagate into the P&L, and the sensitivity analysis in Chapter 7 should be read as an acknowledgement of that dependency rather than as a promise of precision.

The competitive landscape is a snapshot. The CLM, BIM, and point-solution segments are moving quickly — both because incumbent vendors are adding LLM-based capabilities and because new entrants continue to appear — and any comparison table captured in print is subject to being overtaken. The differential contribution identified in Section 3.4 is therefore expressed in structural rather than feature-list terms, but sustained maintenance of that differential is itself an ongoing execution requirement.

Finally, the analysis focuses on Spain as the beachhead market. The argument for geographic expansion sketched in Section 4.5 — Portugal, Italy, Germany, Poland — is commercially plausible and is partially supported by the fact that Spanish IPPs already operate in those markets, but each country brings its own regulatory, linguistic, and procurement nuances that would require country-specific research before the platform could be offered outside Spain at scale.

8.4 Future work and roadmap

ConCoord starts the next 90 days with a working platform, a mapped market, defined pricing, and direct C-suite access inside the Spanish IPP sector. What is missing is commercial validation. The plan for the following quarter is therefore deliberately modest: land one paid pilot, sign a second, and document what worked so the motion can be repeated. The goal is not scale, breadth, or certification — it is evidence.

To reach that evidence, four workstreams advance in parallel rather than in sequence. Product matures from prototype to pilot-ready. Commercial converts warm relationships into signed agreements. Legal and compliance put the minimum paperwork in place to let an infrastructure customer sign without friction. Operations keeps the cadence and closes the first two hires. None of the four is ambitious in isolation; together they cover the ground needed to exit the quarter with a real customer, a real case study, and a plan that no longer depends on the founder's personal network to keep running.

Beyond the first 90 days, two expansion vectors shape the medium-term roadmap. Geographically, the natural extension is into other European renewable markets (Portugal, Italy, Germany, Poland) where many Spanish IPPs already operate. Sectorally, data centres are the most immediate adjacency — they involve construction milestones, SLAs, supplier agreements, grid-connection obligations, and long-term customer contracts, all of which must be tracked against physical reality. Conventional energy, transportation infrastructure, and water/utilities follow the same logic.

An outcome-based pricing component — 5–8% of successfully enforced claims, warranty recoveries, or liquidated damages collected — becomes feasible once 15–20 clients have been operating with ConCoord long enough to produce auditable claim-outcome data. Introducing such a component aligns incentives across vendor and client and is a genuine differentiator in the pitch.

Chapter 9. BIBLIOGRAPHY

[1] UNEF (2024). Informe Anual 2024 – El tiempo de la energía solar. Madrid: Unión Española Fotovoltaica.

[2] APPA Renovables (2024). Estudio del Impacto Macroeconómico de las Energías Renovables en España 2024. Madrid: Asociación de Empresas de Energías Renovables.

[3] Arcadis (2025). 15th Annual Global Construction Disputes Report. Amsterdam: Arcadis NV.

[4] HKA (2025). CRUX Insight – 8th Annual Report. London: HKA Global.

[5] Burford Capital (2024). 2024 Industry Research: Energy Sector. New York: Burford Capital.

[6] OMIE (2024). Reglas de Funcionamiento de los Mercados Diario e Intradiario de Producción de Energía Eléctrica (MIBEL). Madrid: Operador del Mercado Ibérico de Energía.

[7] Iberdrola SA (2024). Annual Report 2023. Bilbao: Iberdrola.

[8] Reuters (2024a). "Saeta Yield secures €1.74 billion debt refinancing", Reuters, 15 March 2024.

[9] TPG (2022). TPG Rise Climate Completes Acquisition of Matrix Renewables. Press release, 1 March 2022. New York: TPG Inc.

[10] Grenergy Renovables (2024). Informe Anual Integrado 2023. Madrid: Grenergy Renovables.

[11] Qualitas Energy (2024). Company Fact Sheet 2024. Madrid: Qualitas Energy.

[12] Reuters (2024b). "Solaria closes €1.7 billion green debt package", Reuters, 2024.

[13] Chambers & Partners (2025). Private Equity 2025 – Spain. London: Chambers & Partners Ltd.

[14] IEA (2025). Renewables 2025 – Analysis and Forecast to 2030. Paris: International Energy Agency.

[15] IRENA (2024). Renewable Capacity Statistics 2024. Abu Dhabi: International Renewable Energy Agency.

[16] Red Eléctrica de España (2024). Informe del Sistema Eléctrico Español 2023. Madrid: Redeia Corporación.

[17] SolarPower Europe (2024). EU Market Outlook for Solar Power 2024–2028. Brussels: SolarPower Europe.

[18] European Parliament and Council of the EU (2018). Directive (EU) 2018/2001 of 11 December 2018 on the promotion of the use of energy from renewable sources (recast), Official Journal of the European Union, L 328/82.

[19] Gobierno de España (2020). Real Decreto 960/2020, de 3 de noviembre, por el que se regula el régimen económico de energías renovables, BOE núm. 289, 4 November 2020.

[20] FIDIC (2017a). Conditions of Contract for EPC/Turnkey Projects – Silver Book, 2nd edn. Geneva: International Federation of Consulting Engineers.

[21] FIDIC (2017b). Conditions of Contract for Construction – Red Book, 2nd edn. Geneva: International Federation of Consulting Engineers.

[22] International Chamber of Commerce (2021). ICC Arbitration Rules. Paris: ICC.

[23] CMAA (2024). Construction Management & Project Management Fee Survey 2024. McLean, VA: Construction Management Association of America.

[24] Clearstory (2024). Clearstory Company Overview and Customer Case Studies. San Francisco, CA: Clearstory.

[25] Crunchbase (2026). Clearstory – Funding Rounds. Available at: https://www.crunchbase.com/organization/clearstory [Accessed 20 April 2026].

[26] Future Market Insights (2024). Contract Lifecycle Management Market Forecast, 2026–2036. Newark, DE: Future Market Insights.

[27] Park, J. S., O'Brien, J., Cai, C. J., Morris, M. R., Liang, P. and Bernstein, M. S. (2023). "Generative Agents: Interactive Simulacra of Human Behavior", Proceedings of UIST 2023.

[28] Wang, L., Ma, C., Feng, X. et al. (2024). "A Survey on Large Language Model based Autonomous Agents", Frontiers of Computer Science, 18(6), 186345.

[29] Hendrycks, D., Burns, C., Chen, A. and Ball, S. (2021). "CUAD: An Expert-Annotated NLP Dataset for Legal Contract Review", Proceedings of NeurIPS 2021, Datasets and Benchmarks Track.

[30] Chalkidis, I., Fergadiotis, M., Malakasiotis, P., Aletras, N. and Androutsopoulos, I. (2022). "LexGLUE: A Benchmark Dataset for Legal Language Understanding in English", Proceedings of ACL 2022, pp. 4310–4330.

[31] Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W., Rocktäschel, T., Riedel, S. and Kiela, D. (2020). "Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks", Proceedings of NeurIPS 2020, pp. 9459–9474.

[32] Szabo, N. (1997). "Formalizing and Securing Relationships on Public Networks", First Monday, 2(9).

[33] Summize (2026). Summize Closes $50M Series B to Scale Contract Intelligence. Press release, January 2026.

[34] Arcadis (2020). Global Construction Disputes Report 2020. Amsterdam: Arcadis NV.

ANEXO I: ALINEACIÓN DEL PROYECTO CON LOS ODS

La Universidad Pontificia Comillas exige, para los Trabajos Fin de Grado, una reflexión explícita sobre la contribución del proyecto a los Objetivos de Desarrollo Sostenible (ODS) de Naciones Unidas. ConCoord, como plataforma de gestión contractual orientada al sector renovable, tiene una relación directa con cuatro ODS — 7, 9, 12 y 13 — que se desarrollan a continuación. La contribución no es incidental: el valor que el producto libera se materializa precisamente a través de una aceleración del despliegue de activos renovables, una mejora de la eficiencia industrial en la operación de infraestructuras críticas y una reducción de la fricción económica que hoy ralentiza la transición energética.

ODS 7: Energía asequible y no contaminante

El ODS 7 persigue garantizar el acceso a una energía asequible, fiable, sostenible y moderna para todos, lo que incluye aumentar sustancialmente la cuota de energías renovables en la matriz global. ConCoord contribuye a esta meta por una vía indirecta pero medible: al reducir la fracción de valor económico que se pierde hoy en los contratos de construcción, O&M y suministro de activos renovables, rebaja el coste total de propiedad de los proyectos y, por tanto, su coste ponderado de capital. Un menor coste de capital se traduce directamente en proyectos que antes no alcanzaban el umbral de rentabilidad y que ahora sí lo hacen, ampliando el volumen de capacidad renovable que el mercado español está dispuesto a financiar sin subvenciones adicionales. La plataforma también acelera la resolución de disputas y el cobro de garantías, lo que mejora el perfil de riesgo que los financiadores asignan al sector y libera liquidez hacia nuevas inversiones.

ODS 9: Industria, innovación e infraestructura

El ODS 9 aboga por construir infraestructuras resilientes, promover la industrialización inclusiva y sostenible y fomentar la innovación. ConCoord es una aplicación concreta de inteligencia artificial agéntica al dominio de la administración contractual de infraestructuras, un ámbito tradicionalmente desatendido por la innovación digital pese a ser determinante en la rentabilidad y resiliencia de los activos. Al trasladar a código auditable una función que hoy depende de procesos manuales, hojas de cálculo y revisiones reactivas, el proyecto contribuye a la modernización tecnológica del tejido industrial español — en particular del segmento de productores independientes de energía renovable, que se sitúa en una posición intermedia entre las grandes utilities y los operadores de pequeña escala y que concentra buena parte del crecimiento esperado en los próximos años. La arquitectura multi-tenant, desplegada sobre infraestructura en la nube con residencia europea, es además coherente con los requisitos de seguridad que NIS2 y el resto del marco regulatorio de infraestructuras críticas están imponiendo al sector.

ODS 12: Producción y consumo responsables

El ODS 12 se centra en garantizar modalidades de consumo y producción sostenibles, lo que incluye una gestión más eficiente de los recursos naturales y una reducción de pérdidas a lo largo del ciclo de vida de los activos. La aportación de ConCoord a este objetivo es específica: mediante la vigilancia continua de obligaciones de disponibilidad, mantenimiento preventivo y garantías de rendimiento, la plataforma permite que los activos operen más cerca de su óptimo técnico y que los incumplimientos contractuales se detecten antes de convertirse en degradación operativa. Esto reduce el desperdicio de energía producida (por indisponibilidad evitable) y alarga la vida útil económica de los equipos, con impactos directos en el consumo de materiales a lo largo del ciclo de reposición. La trazabilidad completa de cada obligación y su evidencia es también la base sobre la que construir reportes ESG verificables, un requisito creciente tanto para inversores institucionales como para la Taxonomía europea.

ODS 13: Acción por el clima

El ODS 13 exige adoptar medidas urgentes para combatir el cambio climático y sus efectos. El despliegue acelerado de capacidad renovable es, en términos cuantitativos, una de las palancas más efectivas para reducir las emisiones de CO_2 asociadas a la generación eléctrica. ConCoord se inserta en esa cadena de valor eliminando una fricción específica — la pérdida de valor en la capa contractual — que hoy actúa como impuesto implícito sobre la velocidad de despliegue. Cada punto de mejora en la eficiencia contractual se traduce en capacidad renovable adicional puesta en operación antes, en garantías ejecutadas que financian nueva inversión y en contratos de largo plazo (PPAs) que se estructuran con menos margen defensivo porque el riesgo residual de ejecución contractual es menor. La contribución de ConCoord a la acción climática es, por tanto, la de un habilitador de la transición: no sustituye a la generación renovable, pero hace que cada euro invertido en ella rinda más en términos de energía limpia efectivamente producida.

ANEXO II: MATERIAL SUPLEMENTARIO

Pricing model spreadsheet

Full pricing and unit-economics calculations are maintained in ConCoord_Pricing_Model.xlsx. The spreadsheet mirrors the assumptions in §7.1 and allows sensitivity analysis across LLM-cost decline, plant count, and pricing-tier mix.

Extended product and architecture description

A longer-form product and architecture description — including additional figures and implementation details — is maintained in Product_and_Architecture_Description.docx as a companion document to this TFG.

Three-month deployment plan

The operational rollout plan for the first commercial quarter is documented in ConCoord_3_Month_Deployment_Plan.docx, which expands on the milestones summarised in §7.4.