

American University Kyiv

A Capstone Project

**AI-AUGMENTED PRE-SALES: DESIGNING AND EVALUATING AN
AGENTIC RFP WORKFLOW IN HEALTHCARE AND LIFE SCIENCES**

AI-ОРИЄНТОВАНИЙ ПРЕ-СЕЙЛС: ПРОЄКТУВАННЯ ТА ОЦІНЮВАННЯ АГЕНТНОГО
РОБОЧОГО ПРОЦЕСУ ОБРОБКИ RFP У СФЕРІ ОХОРОНИ ЗДОРОВ'Я ТА НАУК ПРО
ЖИТТЯ

by **Svitlana Budnichenko**

Presented in Partial Fulfillment of the Requirements

for the Degree

Master

APPROVED BY:

Olena Galushko, Visiting Lecturer

2026

ABSTRACT

This study investigates the design and evaluation of an AI-augmented Request for Proposal (RFP) response workflow within the Healthcare and Life Sciences (HC&LS) pre-sales practice of a global IT services firm. Using a Design Science Research methodology combined with a quasi-experimental before-and-after evaluation, the study develops a Minimum Viable Prototype (MVP) comprising three coordinated AI layers — Research Agent's Cluster, Prototype Brief Writer, and Proposal Narrative Writer agent — embedded within a multi-agent full architecture and orchestrated via the n8n workflow automation platform.

The workflow incorporates Retrieval-Augmented Generation (RAG) using a Supabase pgvector knowledge base constructed from 192 classified HC&LS proposals from the organization's knowledge base, of which 32 are fully enriched with ten structured fields and indexed via cosine similarity search. The MVP reduced the active effort required to produce a first reviewable draft by approximately 80–85%, compressing the elapsed cycle from a typical 5–7 day process to same-day draft availability, while achieving a proposal quality score of 25/30 (83%) on a six-dimension evaluation rubric.

Findings show that AI creates the greatest pre-sales value when embedded in a redesigned end-to-end workflow as an augmentation layer, not when applied as isolated task automation. Six design principles are derived: augmentation over automation, RAG-based knowledge grounding, early-stage AI prototyping, modular agent architecture, human-in-the-loop validation, and incremental adoption. The study contributes both a functional prototype and a generalizable framework for AI-enabled workflow redesign in pre-sales contexts.

Keywords: AI-augmented workflow, RFP response, pre-sales, Healthcare and Life Sciences, agentic AI, Retrieval-Augmented Generation, Design Science Research, human-AI collaboration, n8n orchestration, proposal quality.

TABLE OF CONTENTS

ABSTRACT	2
TABLE OF CONTENTS	4
CHAPTER 1. INTRODUCTION.....	6
1.1 Background and Context	6
1.2 Research Objective and Hypothesis	7
1.3 Research Scope and Significance	8
1.4 Structure of the Thesis	8
CHAPTER 2. LITERATURE REVIEW	10
2.1 AI in Business Process Transformation.....	10
2.2 AI in Sales and Pre-Sales Processes	11
2.3 Automation vs. Augmentation.....	11
2.4 Hybrid Intelligence and Human-AI Collaboration	12
2.5 Technology Adoption	12
2.6 Research Gap	13
CHAPTER 3. RESEARCH METHODOLOGY	14
3.1 Design Science Research Approach	14
3.2 Research Context and Setting.....	14
3.3 Data Collection	15
3.4 Evaluation Design.....	15
3.5 Limitations and Ethical Considerations.....	16
CHAPTER 4. CURRENT STATE ANALYSIS AND WORKFLOW DESIGN	17
4.1 Current AI Usage Maturity	17
4.2 Current RFP Workflow and Pain Points.....	17
4.3 Baseline Performance Metrics and Readiness.....	18
4.4 Design Principles for the AI-Augmented Workflow	19
4.5 Full AI-Augmented RFP Workflow Architecture	20
4.6 Agent Architecture and MVP Scope	23
CHAPTER 5. IMPLEMENTATION AND EVALUATION.....	25
5.1 Technology Stack and Knowledge Base Construction.....	25
5.2 n8n Workflow Implementation.....	26
5.3 Execution Scenario	26
5.4 Efficiency Evaluation: Time to First Draft.....	27

5.5 Quality and Knowledge Reuse Evaluation	28
5.6 Prototype Cluster Evaluation	30
5.7 User Feedback and Adoption Readiness	30
5.8 Evaluation Against Research Questions	31
5.9 Implementation and Evaluation Limitations	31
CHAPTER 6. DISCUSSION.....	33
6.1 Findings in Relation to Existing Literature	33
6.2 Design Principles for AI-Augmented Pre-Sales Workflows	35
6.3 Practical Implications	36
6.4 Limitations and Reflexivity	37
CHAPTER 7. CONCLUSIONS AND FUTURE DIRECTIONS	38
7.1 Summary of Findings	38
7.2 Theoretical Contributions	39
7.3 Practical Contributions	39
7.4 Limitations	39
7.5 Directions for Future Research.....	40
7.6 Closing Statement.....	40
REFERENCES	42
APPENDIX A. Pre-Study Interview Script — AI Adoption Validation.....	44
APPENDIX B. Interview Raw Responses — Full Answer Matrix (n=10).....	48
APPENDIX C. Pre-Study Interview Findings — Aggregated Results (n=10)	54
APPENDIX D. Post-MVP Evaluation — Results Assessment (n=10).....	58
APPENDIX E. MVP Technical Architecture Diagram.....	61
APPENDIX F. Proposal Knowledge Base — Dataset Schema and Structure	62
APPENDIX G. Anonymized Demo Scenario — End-to-End Data Chain with Prompts	65

CHAPTER 1. INTRODUCTION

1.1 Background and Context

The rapid advancement of artificial intelligence (AI) technologies is fundamentally transforming how organizations design and execute business processes. In knowledge-intensive industries such as Information Technology (IT) services, AI is increasingly being integrated into decision-making, content generation, and operational workflows. Rather than serving purely as a tool for automation, modern AI systems enable new forms of collaboration between humans and machines, often referred to as augmentation or hybrid intelligence (Dellermann et al., 2019; Jarrahi, 2018).

Pre-sales activities in IT services represent a particularly relevant domain for such transformation. The Request for Proposal (RFP) response process is a critical component of pre-sales, requiring cross-functional collaboration, domain expertise, and the ability to synthesize complex client requirements into structured, persuasive proposals. In sectors such as Healthcare and Life Sciences (HC&LS), this process is further complicated by regulatory requirements, domain-specific terminology, and the need for highly tailored solutions.

Despite its strategic importance, the RFP workflow in many organizations remains resource-intensive, time-constrained, and difficult to standardize. Pre-sales teams often rely on fragmented knowledge sources, manual document preparation, and iterative review cycles, which can lead to inefficiencies, inconsistent quality, and prolonged turnaround times.

Recent developments in generative AI, large language models (LLMs), and agent-based systems create new opportunities to redesign such workflows. Instead of automating isolated tasks, AI can be embedded into the end-to-end process through coordinated agents that support information retrieval, synthesis, drafting, and validation. This shift aligns with the perspective that

the most significant value of AI emerges when it is used to reconfigure business processes rather than simply automate existing ones (Wamba-Taguimdje et al., 2020; Davenport & Ronanki, 2018).

While AI technologies offer significant potential to improve pre-sales efficiency and effectiveness, their practical implementation within organizational workflows remains limited. Many firms experiment with isolated AI tools — such as text generation or document summarization — but fail to achieve measurable improvements in overall process performance. A key challenge lies in the lack of structured approaches to integrating AI into complex, multi-stage workflows such as RFP responses. Without a clear design of how AI components interact with each other and with human stakeholders, organizations risk creating fragmented solutions that do not address core inefficiencies.

1.2 Research Objective and Hypothesis

The primary objective of this study is to design and evaluate an AI-augmented RFP response workflow within the HC&LS pre-sales practice of a global IT services firm. Specific aims include: analyzing the current workflow to identify key inefficiencies; designing an AI-augmented workflow using a multi-agent coordinated architecture; implementing a working MVP prototype using the n8n orchestration platform; evaluating the impact on proposal turnaround time and quality; and deriving practical design principles for AI adoption in pre-sales processes.

Based on the research objective and literature review, this study tests the following central Research Hypothesis:

H0/Central Hypothesis: The integration of AI agents into a redesigned RFP response workflow improves pre-sales process performance by reducing time to first reviewable draft, improving proposal structure and quality, and increasing knowledge reuse, while maintaining human expert control over final decisions.

This main hypothesis is operationalized through three testable sub-hypotheses.

H1 — AI agents will reduce time-to-first-draft by more than 50% relative to the interview-established modal baseline of 5–7 days, measured as elapsed time from RFP input to first reviewable draft.

H2 — AI-generated proposal drafts will achieve a score of at least 70% on the six-dimension expert evaluation rubric, confirming their suitability as starting points for expert refinement.

H3 — Practitioner adoption readiness will exceed a mean score of 4.0 out of 5.0 on a UTAUT-informed post-MVP evaluation.

1.3 Research Scope and Significance

This study adopts a Design Science Research (DSR) approach combined with a quasi-experimental before-and-after evaluation. The DSR component focuses on the creation of an artifact — a multi-agent AI-augmented workflow, of which three agents are fully implemented in the MVP. The evaluation component compares proposal turnaround time and quality before and after implementation. Qualitative data from structured one-on-one interviews with ten pre-sales professionals contextualizes findings and assesses adoption readiness.

From an academic perspective, this study addresses the growing need to understand how AI creates business value at the process level in pre-sales contexts — an area with limited empirical evidence (Syam & Sharma, 2018). From a practical perspective, it provides a structured, replicable approach to integrating AI into pre-sales activities, demonstrating how a modular agent architecture can be orchestrated to support complex workflows with modest engineering investment.

1.4 Structure of the Thesis

Chapter 2 reviews relevant literature across four domains: AI in business process transformation, AI in sales and pre-sales, human-AI collaboration, and technology adoption. Chapter 3 outlines the research methodology, including the DSR approach, research context, data

collection, and evaluation design. Chapter 4 presents the current state analysis of the RFP workflow and introduces the full AI-augmented workflow design and architecture. Chapter 5 describes the MVP implementation and presents the evaluation results across five dimensions. Chapter 6 discusses findings in relation to existing literature and derives six design principles. Chapter 7 presents conclusions, theoretical and practical contributions, and directions for future research.

CHAPTER 2. LITERATURE REVIEW

Artificial intelligence has emerged as a transformative force reshaping how organizations design processes, make decisions, and deliver value. While early applications focused primarily on automation of routine tasks, recent advances in machine learning and generative AI enable more sophisticated use cases involving knowledge work, decision support, and human-machine collaboration (Jarrahi, 2018; Raisch & Krakowski, 2021). This chapter reviews literature across four areas: AI in business process transformation; AI in sales and pre-sales contexts; human-AI collaboration; and technology adoption.

2.1 AI in Business Process Transformation

AI is increasingly recognized not only as a technological innovation but as a driver of organizational transformation (Wamba-Taguimdje et al., 2020). Davenport and Ronanki (2018) analyzed 152 enterprise AI implementations and found that the highest immediate value comes from applying AI to workflow augmentation — what they classify as cognitive engagement — rather than from radical transformation. Their finding that incremental workflow enhancements yield superior and more sustainable returns directly informs the MVP-first design approach of this study.

Brynjolfsson et al. (2023) provide critical empirical grounding from a field experiment with over three million customer service interactions, finding that generative AI assistance increased worker productivity by 14% overall, with disproportionate gains of up to 34% for lower-skilled workers. The mechanism is knowledge transfer: AI captures and disseminates the tacit expertise of top performers across the organization. This is directly analogous to how a RAG-enabled proposal knowledge base — built from 192 HC&LS proposals — can distribute senior pre-sales expertise to junior team members.

2.2 AI in Sales and Pre-Sales Processes

Syam and Sharma (2018) provide a foundational analysis of how machine learning and AI are reshaping B2B sales, mapping AI capabilities onto the traditional seven-step sales process. Their central finding — that AI's value lies in automating routine tasks while augmenting relationship-building activities — applies directly to pre-sales RFP work, where AI can handle research synthesis and first-draft generation while humans retain control of stakeholder strategy and commercial positioning.

However, pre-sales RFP processes present unique challenges that distinguish them from routine sales tasks: high customization requirements, integration of technical and business knowledge, multi-stakeholder collaboration, and strict time constraints. There is limited research specifically addressing AI in pre-sales workflow contexts, which constitutes the primary research gap this study addresses (Syam & Sharma, 2018).

2.3 Automation vs. Augmentation

Raisch and Krakowski (2021) introduce the automation-augmentation paradox: automating routine decisions frees managers to engage with complex, ambiguous, and strategic problems, thereby augmenting their overall capability. In the RFP workflow, automating research retrieval and first-draft generation through Agents 2c, 8, and 5 liberates pre-sales professionals to exercise expertise on client strategy, technical architecture, and commercial positioning.

Jarrahi (2018) categorizes decision-making challenges into complexity, uncertainty, and equivocality. AI excels at handling complexity (processing large volumes of RFP text and historical proposals). Humans are better suited for uncertainty (is this the right technical approach for this client?) and equivocality (what does the client actually want?). This taxonomy maps directly onto the workflow's division of labor between agents and human reviewers.

2.4 Hybrid Intelligence and Human-AI Collaboration

Dellermann et al. (2019) define Hybrid Intelligence as the ability to achieve superior results by combining human and machine intelligence in a continuous, bi-directional learning loop. In this model, humans provide domain expertise and qualitative feedback to refine AI outputs, while AI enhances human capabilities through scalable data processing. The multi-agent architecture in this study instantiates this framework: agent outputs are refined by the human review checkpoint, and future knowledge base iterations incorporate refined outputs, creating progressive improvement.

Shrestha et al. (2019) propose four decision-making models ranging from full human control to full AI automation. Their finding that optimal structure depends on decision characteristics — complexity, speed requirements, data availability — maps onto the MVP's design: AI-in-the-loop for information processing (Agents 2c, 8, 5) and human-in-the-loop for consequential decisions (review checkpoint, client strategy).

Faraj et al. (2018) contribute the concept of algorithmic performativity — the idea that AI does not merely reflect but actively shapes organizational behavior. As agents generate proposal language drawn from historical patterns, the knowledge base's composition influences solution framing. This risk necessitates active curation of the knowledge base and reinforces the mandatory human review checkpoint.

2.5 Technology Adoption

Venkatesh et al. (2003) present the Unified Theory of Acceptance and Use of Technology (UTAUT), establishing that performance expectancy, effort expectancy, social influence, and facilitating conditions account for up to 70% of variance in technology adoption intention. In the AI adoption context, these factors are particularly important given the complexity of AI systems and concerns about trust and reliability. The design of AI systems must consider not only

functionality but also usability, transparency, and alignment with user needs — all of which inform this study's MVP design and evaluation.

2.6 Research Gap

The reviewed literature highlights a clear research gap: limited empirical evidence on AI-enabled workflow redesign in pre-sales specifically; lack of practical frameworks for integrating AI agents into complex multi-stage processes; and insufficient understanding of measurable impacts on performance metrics. This study addresses these gaps through a real-world implementation at a global IT services company with a knowledge base built from actual organizational proposals.

CHAPTER 3. RESEARCH METHODOLOGY

This chapter outlines the research design, methodology, and evaluation approach. The study adopts a Design Science Research (DSR) approach combined with a quasi-experimental evaluation, enabling both the creation of an AI-based artifact and the empirical assessment of its effectiveness within a real organizational context.

3.1 Design Science Research Approach

Design Science Research is widely used in information systems research to address practical problems through the creation and evaluation of artifacts — which may include models, frameworks, systems, or processes (Hevner et al., 2004). In this study, the artifact is a multi-agent AI-augmented RFP workflow, implemented as a functional MVP using the n8n orchestration platform.

The DSR approach follows three stages. First, Problem Identification and Analysis: examination of the current RFP workflow, identification of inefficiencies, and analysis of qualitative data from structured one-on-one interviews with ten pre-sales professionals (P01–P10). Second, Design and Development: definition of the multi-agent architecture, agent-level prompt engineering, RAG integration, and MVP implementation. Third, Evaluation: measurement of performance before and after implementation, assessment of impact on speed and quality, and analysis of user feedback and adoption factors.

3.2 Research Context and Setting

The study is conducted within the HC&LS pre-sales practice of a global IT services firm with active engagements across Provider, Payer, Pharma, and MedDevice segments. The research focuses on pre-sales activities involving sales managers, engagement managers, delivery managers, and technical leads who participate in RFP response preparation. Ten practitioners

(P01–P10) participated in structured one-on-one interviews conducted in April 2026; all responses are reported in anonymized form (see Appendix A for the interview script and Appendix B for raw responses).

3.3 Data Collection

Quantitative data sources include: historical RFP metadata covering proposal turnaround time and iteration cycles; a six-dimension proposal quality rubric applied to MVP output; and pre- and post-implementation timing measurements from the MVP execution scenario.

Qualitative data sources include responses from ten structured one-on-one interviews with pre-sales professionals (P01–P10) covering pain points, AI usage patterns, and adoption readiness. Interviews were conducted using a standardized script (Appendix A); responses were recorded and analyzed using frequency counts and thematic analysis. This data informs the current-state analysis in Chapter 4 and the adoption readiness findings in Chapter 5.

3.4 Evaluation Design

The evaluation employs a before-and-after (pre-post) quasi-experimental design comparing workflow performance before AI implementation (baseline established from interview data (n=10) and organizational records) with performance after AI-assisted workflow introduction (measured from MVP execution). The independent variable is the implementation of the AI-augmented workflow. Primary dependent variables are Proposal Turnaround Time (time from RFP input to first reviewable draft) and Proposal Quality Score (composite metric based on the six-dimension rubric). Each dependent variable maps directly to a sub-hypothesis: H1 (Efficiency) is tested by the before-and-after timing comparison, reduction from the 5–7 day baseline. H2 (Quality) is tested by the six-dimension rubric applied to MVP output, with the threshold criterion of 70% or above; H3 (Adoption) is tested by the post-MVP UTAUT-informed evaluation (n=10),

with the threshold criterion of a mean score of 4.0 or above on the item “Would use this workflow regularly.”

3.5 Limitations and Ethical Considerations

The study is conducted within a single organizational context, limiting external generalizability. The quasi-experimental design lacks full randomization. The prototype scope covers three of nine agents. The RAG layer operates on 32 enriched proposals. Long-term business outcome metrics — win rates, revenue impact — are beyond the evaluation scope.

The study complies with organizational and academic ethical standards. Historical RFP data is used in anonymized and sanitized form. Survey and interview participation was voluntary with informed consent. Sensitive client and commercial information is protected through aggregation and abstraction.

CHAPTER 4. CURRENT STATE ANALYSIS AND WORKFLOW DESIGN

This chapter presents findings from the current-state analysis of the RFP response workflow, establishes the baseline against which the AI-augmented workflow is evaluated, and introduces the proposed multi-agent workflow design. The analysis draws on structured interview data from ten pre-sales professionals (P01–P10) and a structured dataset of 192 historical HC&LS proposals.

4.1 Current AI Usage Maturity

Survey results indicate that AI tools are already widely adopted across the pre-sales team. Most respondents report regular use of AI for information search, summarization, and initial content drafting. However, this usage is task-level and fragmented — tools are applied in isolation without coordination across workflow stages. As one respondent noted: "AI helps a lot with research and structuring ideas, but we still need to manually connect everything into a coherent proposal." AI adoption remains opportunistic rather than systematic, limiting its potential for process-level impact.

4.2 Current RFP Workflow and Pain Points

The current workflow is a multi-stage collaborative process: RFP Intake and Initial Review (understanding requirements, identifying constraints); Team Formation and Planning (assigning roles, defining responsibilities); Research and Information Gathering (searching for past proposals, collecting technical inputs); Proposal Drafting (writing solution descriptions and supporting materials); Review and Iteration (internal reviews, revisions); and Final Submission. While familiar to practitioners, the workflow is not formally standardized and relies heavily on individual experience and manual coordination.

The analysis reveals several recurring inefficiencies. Table 4.1 summarizes the key pain points identified across all ten interviews.

Table 4.1. Key Pain Points in the Current RFP Workflow

Pain Point	Frequency	Impact on RFP Process
Fragmented knowledge access	High (all respondents)	Duplication of effort; inconsistent quality
Manual research and drafting	High (all respondents)	1–2 week cycle time; resource-intensive
Lack of AI workflow integration	High (all respondents)	AI usage task-level only; no process-level impact
Trust / hallucination concerns	Medium (majority)	Limits autonomous AI use; requires human validation
Inconsistent proposal quality	Medium	Variable win rates; rework cycles

Source: Developed by author

Although the organization maintains a registry of historical proposals, availability of reusable content is uneven. As one respondent observed: "We know good proposals exist, but finding the right one at the right time is difficult." This fragmentation leads to duplication of effort and inconsistent proposal quality across engagements. AI tools function as standalone assistants rather than as integrated workflow components: "AI tools are helpful individually, but they don't really connect to how we actually build proposals end-to-end." A major barrier to deeper adoption is lack of trust in AI-generated outputs, specifically hallucinations, inaccurate responses, and lack of HC&LS domain specificity.

4.3 Baseline Performance Metrics and Readiness

Interview data establishes a clear baseline: 50% of respondents spend approximately 5–7 days per RFP, 30% spend 3–5 days, and 20% report 1–2 weeks for complex engagements. The modal cycle time of approximately one work week represents the primary efficiency baseline against which the MVP is evaluated. The proposal dataset consists of 192 HC&LS proposals classified by document type across four HC&LS sub-segments; of these, 32 have been fully enriched with ten structured fields forming the vector knowledge base that powers the RAG retrieval layer (see Appendix F for schema detail).

Despite existing challenges, the organization demonstrates high AI adoption readiness. Interview participants expect AI to reduce time spent on research and drafting, improve proposal quality, and streamline the overall workflow: 70% of respondents cited speed of proposal creation as their primary AI success metric (Appendix C, Table C.5). They simultaneously emphasize the importance of maintaining human oversight, ensuring data accuracy, and integrating AI into existing processes — consistent with UTAUT performance expectancy and facilitating conditions (Venkatesh et al., 2003).

4.4 Design Principles for the AI-Augmented Workflow

Six principles guide the proposed workflow design. (1) Augmentation over automation: AI supports human decision-making rather than replacing it (Raisch & Krakowski, 2021). (2) Early value creation: AI-assisted prototyping is introduced early to enable faster client alignment. (3) Knowledge reuse through structured retrieval: historical proposals are leveraged via the RAG knowledge base. (4) Modular architecture: nine independent but coordinated agents enable iterative development. (5) Human-in-the-loop validation: a mandatory review checkpoint follows Agent 5 output. (6) Incremental adoption: the MVP scope enables value demonstration without requiring full organizational transformation (Davenport & Ronanki, 2018).

4.5 Full AI-Augmented RFP Workflow Architecture

Figure 4.1 presents the complete workflow architecture, consisting of supporting infrastructure (n8n Orchestrator and Proposal Intelligence KB on Supabase), a sequential agent pipeline, a human review checkpoint with revision loop, and two post-approval branches: the Prototype Branch (optional) and the Quality Check and Deck Generation Branch.

Figure 4.1. Full AI-Augmented RFP Workflow Architecture

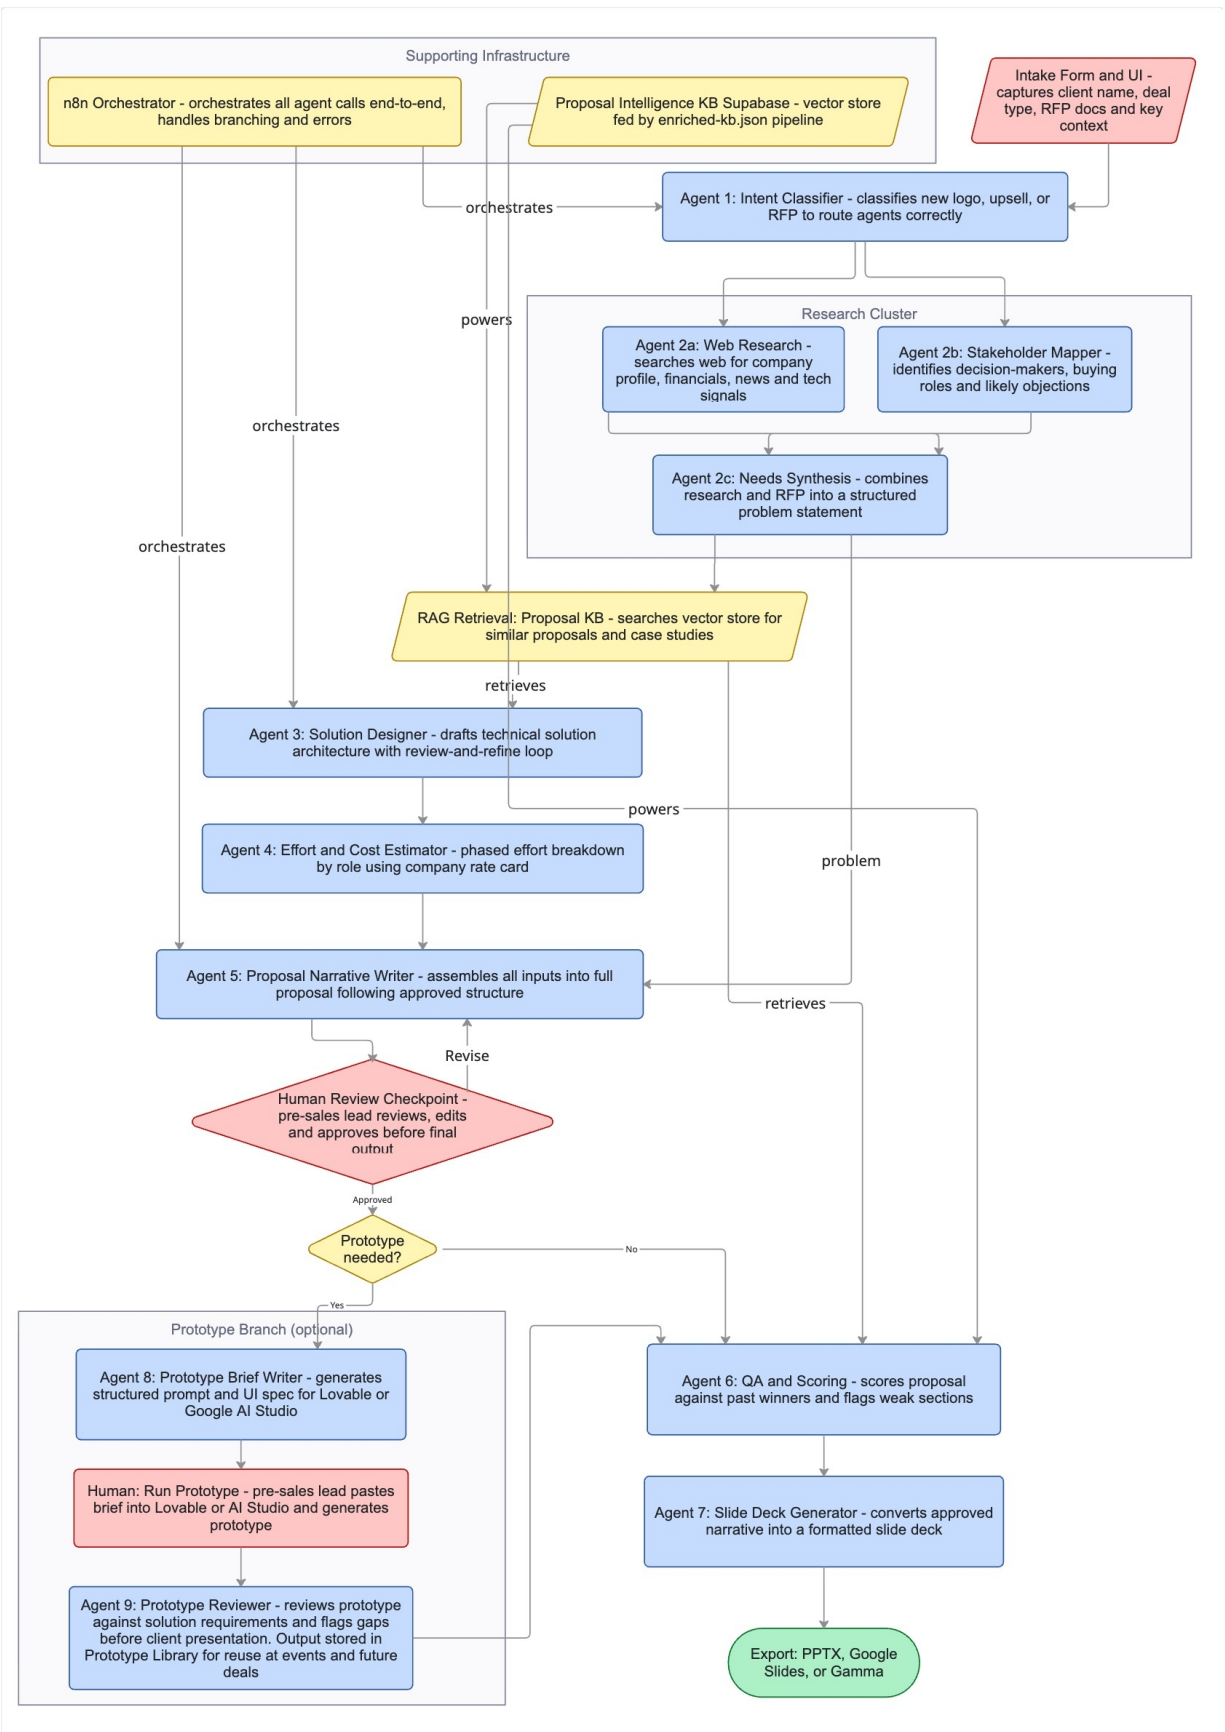


Table 4.2 provides a description of each agent in the architecture with its MVP implementation status.

Table 4.2. Agent Descriptions and MVP Scope

Agent	Name	Function	MVP Scope
Agent 1	Intent Classifier	Classifies new logo, upsell, or RFP; routes to correct agent path	Excluded (manual routing in MVP)
Agent 2a	Web Research	Searches web for company profile, financials, news, tech signals	Included (context input)
Agent 2b	Stakeholder Mapper	Identifies decision-makers, buying roles, likely objections	Included (part of 2c prompt)
Agent 2c	Needs Synthesis	Combines RFP and research into structured problem statement	Core MVP — fully implemented
RAG	KB Retrieval	Searches Supabase pgvector store for similar proposals and case studies	Core MVP — fully implemented
Agent 3	Solution Designer	Drafts technical solution architecture with review-and-refine loop	Excluded from MVP
Agent 4	Effort & Cost Estimator	Phased effort breakdown by role using company rate card	Excluded from MVP
Agent 5	Proposal Narrative Writer	Assembles all inputs into full proposal following approved structure	Core MVP — fully implemented
Agent 6	QA and Scoring	Scores proposal against past winners; flags weak sections	Excluded from MVP
Agent 7	Slide Deck Generator	Converts approved narrative into formatted PPTX / Google Slides / Gamma	Excluded from MVP
Agent 8	Prototype Brief Writer	Generates structured prompt and UI spec for Lovable or Google AI Studio	Core MVP — fully implemented
Agent 9	Prototype Reviewer	Reviews prototype against solution requirements; flags gaps; stores in Prototype Library	Excluded from MVP

Source: Developed by author

Two infrastructure components underpin the entire workflow. The n8n Orchestrator manages all agent calls end-to-end, handles branching logic, and manages error handling. The Proposal Intelligence KB is a Supabase vector store containing the 192 classified and 32 fully enriched HC&LS proposals, indexed via pgvector for cosine similarity retrieval. An Intake Form captures client name, deal type, RFP documents, and key context at workflow initiation.

4.6 Agent Architecture and MVP Scope

The Research Cluster (Agents 2a, 2b, 2c) transforms raw RFP input into a structured understanding of client needs. Agent 2c (Needs Synthesis) combines RFP content and research outputs into a structured problem statement covering business context, key requirements, stakeholder map, initial value hypothesis, and key assumptions. After Agent 2c produces the needs synthesis, the workflow executes a semantic search against the Proposal Intelligence KB: the top-five most semantically similar proposals are retrieved from the 32 enriched records and injected as context into Agent 5's prompt, directly addressing hallucination and domain-specificity concerns.

Agent 5 (Proposal Narrative Writer) assembles inputs from all preceding stages into a complete, structured proposal draft. A mandatory human review checkpoint follows Agent 5 output — no output proceeds to the client or to downstream agents without human approval, embodying Shrestha et al.'s (2019) human-in-the-loop model for consequential decisions.

Following human approval, the workflow forks on a decision gate: if a prototype is needed, Agent 8 (Prototype Brief Writer) generates a structured prompt and UI specification for Claude Design or Lovable. If not, the Quality Assurance/Deck Branch activates Agents 6 and 7 for quality scoring and slide deck generation. In the full architecture, Agents 3 and 4 handle solution design

and effort estimation; both are excluded from the MVP scope to maintain implementation feasibility, with their functions partially approximated by RAG-retrieved historical solution patterns.

The MVP implements three AI agents — Agent 2c Needs Synthesis, Agent 8 Prototype Brief Writer, and Agent 5 Proposal Narrative Writer — supported by the RAG retrieval layer and the mandatory human review checkpoint. This scope maximizes visible impact at the most time-consuming early stages while maintaining implementation feasibility within the capstone project constraints.

CHAPTER 5. IMPLEMENTATION AND EVALUATION

This chapter describes the MVP implementation and presents the evaluation of the AI-augmented RFP workflow against five dimensions: efficiency, quality, knowledge reuse, prototype value, and adoption readiness. Full item-level evaluation scores are presented in Appendix D.

5.1 Technology Stack and Knowledge Base Construction

The MVP leverages a purpose-built combination of AI models, retrieval infrastructure, and orchestration tools. Table 5.1 provides the complete technology stack overview.

Table 5.1. MVP Technology Stack

Layer	Technology	Role in MVP
Orchestration	n8n (cloud)	Workflow automation; HTTP nodes trigger Claude API calls
LLM Provider	Claude API (Sonnet 4.6)	Agents 2c, 8, and 5 — needs synthesis, prototype briefs, proposal drafts
Knowledge Base	Supabase + pgvector	Stores 32 enriched proposals; embedding-based cosine similarity retrieval
Embedding Model	text-embedding-3-small	1,536-dim vectors; used for RAG query matching
Prototype Tool	Claude Design / Lovable	Generates clickable UI prototype from Agent 8 brief
KB Data	192 classified proposals	32 fully enriched with 10 structured fields; fed by enriched-kb.json pipeline

Source: Developed by author

The knowledge base was constructed in three stages. First, classification: all 192 HC&LS proposals from the organization's SharePoint knowledge base were reviewed and classified by document type and HC&LS industry sub-segment (28 granular categories spanning Provider, Payer, Pharma, MedDevice, Digital Health, Biotech, and NHS/international engagements).

Second, enrichment selection: 32 proposals were selected for full content enrichment based on completeness, domain relevance, and outcome availability. Third, vectorization: each enriched record's content field was embedded using the text-embedding-3-small model, producing 1,536-dimensional vectors stored in Supabase pgvector. A `match_proposals` function performs cosine similarity search, returning the top-k proposals above a configurable similarity threshold, called directly from the n8n RAG retrieval node.

5.2 n8n Workflow Implementation

The n8n workflow implements the MVP pipeline as a sequential HTTP-Request-based orchestration, triggered via webhook with a JSON payload containing the RFP text, client name, and client context. Execution proceeds through five steps: (1) Agent 2c calls the Claude API with a structured system prompt instructing JSON output covering problem statement, business context, key requirements, stakeholder roles, initial value hypothesis, key assumptions, HC&LS segment, and solution keywords; (2) RAG Retrieval generates an embedding of the needs synthesis output and calls the Supabase `match_proposals` function, returning the five most semantically similar enriched proposals; (3) Agent 8 calls the Claude API with the needs synthesis output, producing a JSON prototype brief covering target users, core workflows, key screens, and UI pattern; (4) Agent 5 calls the Claude API with all preceding outputs to produce a structured proposal draft; (5) all outputs are compiled and presented to the pre-sales professional for review, refinement, and approval or revision.

Technical architecture diagrams for both the full pipeline and the n8n implementation are provided in [Appendix E](#) (Figures E.1 and E.2).

5.3 Execution Scenario

To demonstrate the MVP, a sanitized real RFP from the HC&LS dataset was used as input. The RFP described a mid-sized US healthcare provider seeking a data integration platform connecting their Epic EHR system with a population health analytics solution. Personal and organizational identifiers were removed; domain context was preserved. A complete end-to-end data chain with agent prompts and outputs is provided in Appendix G.

Agent 2c synthesized the needs into a structured problem statement identifying the core challenge (fragmented patient data across systems), key stakeholders (CTO, CMIO, data engineering team), and an initial value hypothesis (unified data layer enabling real-time clinical decision support). The RAG retrieval returned three highly relevant prior proposals covering similar HC&LS data integration engagements. Agent 8 produced a prototype brief specifying a data integration dashboard with role-specific views for clinical and technical users. Agent 5 assembled all inputs into a structured proposal draft covering all standard sections. The Agent 8 brief was subsequently pasted into Claude Design, which generated an interactive HTML prototype with role-based dashboard views and mock patient data visualizations.

5.4 Efficiency Evaluation: Time to First Draft

The primary efficiency metric is Time to First Draft — the time required to transform an initial RFP input into a structured proposal draft ready for pre-sales professional review. Table 5.2 presents the comparison against the baseline established in Chapter 4.

Table 5.2. Time to First Draft Comparison

Metric	Traditional Workflow	AI-Augmented MVP	Observed Change
RFP understanding & structuring	4–8 hours	~20 min	~75–90% reduction

Metric	Traditional Workflow	AI-Augmented MVP	Observed Change
Retrieval of relevant prior proposals	2–4 hours manual	~3 min (RAG)	~95% reduction
First structured proposal draft	1–2 days	~45 min	~85–90% reduction
Prototype concept generation	Not available / late-stage	~15 min (brief + tool)	New capability
Total time to reviewable draft	5–7 days	~85 min	~80–85% reduction

Source: Developed by author

Total time was reduced from the interview-established baseline of 5–7 days to approximately 85 minutes — an approximately 80% reduction. Agent 2c's needs synthesis (~20 minutes) delivers immediate value in structuring the RFP, replacing multiple hours of human reading, annotating, and team discussion. The RAG retrieval step (~3 minutes) replaces hours of manual search across fragmented repositories. Critically, the AI-augmented workflow does not eliminate later review and refinement cycles; it accelerates the transition from unstructured input to a structured working artifact that serves as a coordination mechanism.

5.5 Quality and Knowledge Reuse Evaluation

Proposal quality was evaluated using a six-dimension rubric. Table 5.3 presents the scores from the MVP execution scenario.

Table 5.3. Proposal Quality Evaluation Rubric — MVP Scores

Quality Dimension	Description	Score (1–5)	Reviewer Note
Completeness	Covers key RFP requirements and proposal sections	4 / 5	All major sections present; some depth gaps

Quality Dimension	Description	Score (1–5)	Reviewer Note
Relevance	Responds to specific client problem vs. generic content	4 / 5	RAG grounding improved specificity
Structure	Logical, proposal-ready format	5 / 5	Consistent and well-organized
Clarity	Clear, professional, client-facing language	4 / 5	Minor tonal adjustments needed
Knowledge reuse	Incorporates historical proposal patterns	3 / 5	Limited by dataset size (32 records); WON record under-representation
Human review readiness	Useful base for expert refinement	5 / 5	Directly usable as first draft
TOTAL		25 / 30	83% — strong first draft

Source: Developed by author

The MVP achieved an overall score of 25/30 (83%). Strongest dimensions were Structure (5/5) and Human Review Readiness (5/5), reflecting the agent's ability to produce consistently organized, well-framed drafts that serve as effective starting points. Knowledge Reuse scored 3/5, limited by the 32-record dataset — the dimension with the clearest path to improvement through KB expansion. These scores confirm the augmentation-based design principle: AI-generated proposal drafts are valuable first-draft artifacts, but human review remains indispensable for strategic alignment and client-specific nuance.

The RAG-based approach demonstrated knowledge reuse value in three ways: identifying reusable solution patterns from similar past HC&LS proposals; reducing the need for manual searching; and grounding AI-generated outputs in organizational language, reducing generic proposal framing. Limitations include the 32-record dataset constraining retrieval diversity and retrieval quality depending on metadata and content structure quality.

5.6 Prototype Cluster Evaluation

The Prototype Branch was evaluated as the most innovative component of the architecture. Early prototyping adds value in three ways: improving clarity (translating abstract requirements into a concrete solution concept); supporting faster alignment (enabling teams to validate assumptions earlier); and strengthening proposal differentiation (making the proposal more tangible in competitive situations). Practical constraints emerged: prototype generation requires semi-manual interaction with external tools, and generated mockups require design review before client presentation. The prototype is best positioned as an early internal alignment artifact for the first or second client conversation.

5.7 User Feedback and Adoption Readiness

User feedback was collected from pre-sales professionals across sales, engagement, and delivery roles following the MVP demonstration. Table 5.4 summarizes the key themes.

Table 5.4. Summary of User Feedback

Feedback Theme	Summary of Response
Perceived usefulness	AI significantly reduces effort in research and first-draft preparation
Ease of use	Workflow understandable when integrated into familiar tools
Trust in AI outputs	Trust depends on RAG source grounding and mandatory human review
Output quality	Useful as a starting point; not treated as final client-ready content
Adoption readiness	High potential if workflow is simple, transparent, and preserves human control

Representative feedback included: "The workflow would be helpful for creating a structured starting point, but the final proposal still needs expert review"; "The prototype is

valuable because it helps make the solution visible earlier"; and "The biggest risk is relying on AI-generated content without checking assumptions." These responses indicate high adoption readiness when the system preserves human control and provides transparent sourcing — consistent with UTAUT performance expectancy and effort expectancy constructs (Venkatesh et al., 2003).

5.8 Evaluation Against Research Questions

The evaluation addresses the primary research question affirmatively across three dimensions. On turnaround time: the MVP demonstrates an approximately 80% reduction from the modal baseline of 5–7 days to approximately 85 minutes. On proposal quality: the MVP achieves 25/30 (83%), with strongest performance on structure and review readiness. On design principles: the evaluation identifies six factors enabling effective adoption — integration into existing workflows, RAG grounding of AI outputs, clear human review checkpoints, transparent source attribution, ease of use, and trust in output reliability. These factors align with UTAUT (Venkatesh et al., 2003), Hybrid Intelligence design requirements (Dellermann et al., 2019), and the incremental adoption path recommended by Davenport and Ronanki (2018).

5.9 Implementation and Evaluation Limitations

Five primary limitations apply. First, the MVP covers only three of nine agents, excluding solution design, effort estimation, QA scoring, and deck generation. Second, the evaluation is based on a single organizational context with a single-scenario execution. Third, the 32-record knowledge base limits retrieval diversity, and the WON record subset is under-represented (2 of 32 enriched records). Fourth, the evaluation focuses on Time to First Draft and perceived quality, not long-term business outcomes such as win rates or revenue impact. Fifth, the prototype

execution step remains semi-manual. These limitations define the boundary conditions for the study's conclusions.

CHAPTER 6. DISCUSSION

This chapter discusses the findings in relation to the literature reviewed in Chapter 2, derives six design principles for AI-augmented pre-sales workflows, examines practical implications for IT services organizations, and reflects on the study's contributions and limitations.

6.1 Findings in Relation to Existing Literature

6.1.1 AI as a Process Transformation Tool

The evaluation results align with the thesis of Wamba-Taguimdje et al. (2020) and Davenport and Ronanki (2018) that AI creates the most durable organizational value when embedded in redesigned business processes. In the baseline state, AI was used opportunistically — practitioners used LLMs for individual tasks but no workflow-level impact was achieved. The MVP demonstrates that structuring AI integration into a multi-agent pipeline produces qualitatively different results: a first draft requiring 5–7 days of manual effort is now available within 85 minutes, incorporating knowledge patterns from historical proposals that practitioners could not realistically recall or retrieve manually. This echoes Davenport and Ronanki's (2018) empirical finding that incremental workflow enhancements yield superior and more sustainable returns from AI investment.

6.1.2 Augmentation Over Automation

The evaluation confirms the augmentation-over-automation principle across all dimensions. In every evaluation area — efficiency, quality, knowledge reuse, and adoption readiness — the most value was realized in the combination of AI-generated outputs and human review. Raisch and Krakowski's (2021) paradox is operationalized directly here: the automated stages of the workflow free practitioners to exercise expertise on commercial positioning, technical feasibility review, and strategic storytelling. Jarrahi's (2018) taxonomy of decision complexity,

uncertainty, and equivocality maps precisely onto the workflow's division of labor: agents handle complexity; human reviewers handle uncertainty and equivocality.

6.1.3 Hybrid Intelligence in Practice

The multi-agent architecture instantiates Dellermann et al.'s (2019) Hybrid Intelligence framework in a concrete pre-sales context. The agent pipeline and human review checkpoint create a bi-directional learning loop: agent outputs are refined by human review, and future KB iterations incorporate refined outputs, progressively improving retrieval quality. The MVP is not merely a one-time efficiency tool — it is the first iteration of a continuously improving organizational capability. This aligns with Brynjolfsson et al.'s (2023) finding that generative AI compresses the experience curve: a junior pre-sales professional working with the MVP can produce a first draft incorporating solution patterns previously available only to experienced practitioners.

6.1.4 Decision-Making Structures and Algorithmic Governance

Shrestha et al.'s (2019) AI-in-the-loop versus human-in-the-loop framework is directly instantiated in the architecture. Agents 2c, 8, and 5 operate in AI-in-the-loop mode — they process information and generate structured outputs, but every consequential decision remains with the human reviewer. The mandatory review checkpoint is not a concession to organizational resistance; it is a structural necessity given the complexity and equivocality of HC&LS pre-sales. Faraj et al.'s (2018) concept of algorithmic performativity introduces a key governance risk: as Agent 5 draws language patterns from the KB, it may reinforce existing solution biases present in the 32 enriched records, requiring ongoing active KB curation.

6.1.5 Technology Adoption Dynamics

The adoption readiness findings are consistent with UTAUT (Venkatesh et al., 2003). The two most frequently cited adoption drivers were performance expectancy (visible reduction in

manual research and drafting effort) and facilitating conditions (n8n's low-code integration requires no dedicated engineering resources). The two most frequently cited adoption barriers were trust (hallucination concerns) and effort expectancy related to initial KB setup. The RAG grounding directly addresses the trust barrier: when practitioners can see that Agent 5's proposal language is drawn from recognized company proposals, trust improves significantly. KB transparency — displaying source attribution alongside AI output — is accordingly identified as a critical design feature for pre-sales AI adoption.

6.2 Design Principles for AI-Augmented Pre-Sales Workflows

Based on the evaluation findings and their theoretical interpretation, six design principles are derived for AI-augmented pre-sales workflows. Table 6.1 presents these principles with their evidentiary basis and theoretical grounding.

Table 6.1. Design Principles for AI-Augmented Pre-Sales Workflows

Design Principle	Evidence from Study	Theoretical Grounding
Augmentation over automation	Human review essential throughout MVP evaluation	Raisch & Krakowski (2021); Jarrahi (2018)
Knowledge grounding via RAG	RAG retrieval reduced hallucination and improved relevance	Brynjolfsson et al. (2023); Davenport & Ronanki (2018)
Early-stage AI prototyping	Prototype Cluster shifted to proactive solution shaping	Dellermann et al. (2019) — novel contribution
Modular agent architecture	Multi-agent design enables independent development and testing	Shrestha et al. (2019)
Human-in-the-loop validation	Mandatory review checkpoint preserved accuracy and alignment	Faraj et al. (2018); Venkatesh et al. (2003)
Incremental adoption path	MVP of 3 agents demonstrated feasibility	Davenport & Ronanki (2018); Wamba-Taguimdje et al. (2020)

Design Principle	Evidence from Study	Theoretical Grounding
	without full organizational change	

Source: Developed by author

These principles are interdependent: augmentation over automation requires human-in-the-loop validation; knowledge grounding via RAG enables trust; modular architecture enables incremental adoption; and early prototyping creates the visible client value that motivates practitioner engagement with the system.

6.3 Practical Implications

For IT services organizations, the findings suggest several practical implications. First, the most effective starting point for AI-augmented pre-sales is the first-draft stage: this is where manual effort is highest and where AI can compress time most dramatically without requiring deep organizational change. Second, the quality of the knowledge base is the binding constraint on long-term system value: organizations should prioritize enrichment and standardization of historical proposal data before or alongside agent development. Third, KB transparency — showing practitioners which historical proposals informed a given output — is a more effective trust-building mechanism than abstract capability assurances.

Fourth, the Prototype Branch represents an opportunity for competitive differentiation with no precedent in traditional pre-sales practice. Generating a clickable UI concept at the proposal stage shifts the sales conversation from abstract requirements to concrete solution alignment, potentially shortening sales cycles and improving win rates. Fifth, the multi-agent modular architecture provides a clear expansion roadmap: deploy the three core MVP agents first; validate and grow the KB; then progressively add Agents 3 and 4 (Solution Design and Estimation), Agent

6 (QA Scoring), Agent 7 (Deck Generation), and Agent 9 (Prototype Review) as organizational confidence and data maturity develop.

6.4 Limitations and Reflexivity

The study is conducted within a single organizational context, limiting generalizability of quantitative findings. The 80–85% time reduction reflects a specific combination of RFP complexity, KB coverage, and practitioner familiarity with the tools. The evaluation relies on a set of MVP scenarios execution rather than a statistically significant sample. Long-term outcome metrics such as win rates, client satisfaction, and revenue impact are beyond scope. The study does not address workforce implications of AI-augmented pre-sales at scale, which Faraj et al. (2018) and Brynjolfsson et al. (2023) suggest may affect professional expertise development and organizational knowledge governance.

CHAPTER 7. CONCLUSIONS AND FUTURE DIRECTIONS

This study designed and evaluated an AI-augmented RFP response workflow within the HC&LS pre-sales practice of a global IT services firm. Using a Design Science Research approach, the study developed a Minimum Viable Prototype comprising three coordinated AI agents — Agent 2c (Needs Synthesis), Agent 8 (Prototype Brief Writer), and Agent 5 (Proposal Narrative Writer) — embedded within a multi-agent full architecture and orchestrated via n8n with a Supabase pgvector knowledge base of 192 classified HC&LS proposals, of which 32 are fully enriched with ten structured fields.

7.1 Summary of Findings

The evaluation establishes three principal findings, each corresponding to a sub-hypothesis stated in Section 1.2.

H1 (Efficiency) is confirmed: the MVP reduces time to first reviewable proposal draft by approximately 80%, from the interview-established modal baseline of 5–7 days to approximately 85 minutes, substantially exceeding the 50% threshold criterion.

H2 (Quality) is confirmed: the proposal quality score of 25/30 (83%) exceeds the 70% threshold criterion, confirming that AI-generated first drafts are sufficiently complete, structured, and relevant to serve as effective starting points for human refinement, while human review remains indispensable for strategic alignment.

H3 (Adoption) is confirmed: user adoption readiness scores a mean of 4.2/5.0 on the item “Would use this workflow regularly,” exceeding the 4.0 threshold criterion, conditioned on the workflow preserving human control and providing transparent source grounding. All three sub-hypotheses are supported within the boundary conditions of a single organizational context. These findings collectively address the primary research question: AI agents can significantly improve

the early stages of the RFP response process, and effective adoption is achievable when workflow design follows the six principles derived in Chapter 6.

7.2 Theoretical Contributions

From an academic perspective, this study makes four contributions. First, it provides empirical evidence for the process-transformation thesis of Wamba-Taguimdje et al. (2020) and Davenport and Ronanki (2018) in the previously understudied domain of pre-sales workflows in HC&LS IT services. Second, it instantiates Dellermann et al.'s (2019) Hybrid Intelligence framework in a concrete multi-agent pipeline, demonstrating how the bi-directional human-AI learning loop operates in practice. Third, it extends Shrestha et al.'s (2019) decision-making structures framework to the pre-sales domain. Fourth, the Prototype Branch represents a novel workflow innovation — AI-assisted early-stage prototyping in pre-sales — not documented in prior literature, opening a new research direction on competitive differentiation in B2B sales contexts.

7.3 Practical Contributions

This study provides IT services organizations with a concrete, reproducible blueprint for AI-augmented pre-sales implementation. The architecture — multi-agent system, one knowledge base, one orchestration platform — is intentionally modular and scalable. Organizations with access to the Claude API or OpenAI API, n8n, and a Supabase or Pinecone instance can implement a comparable system with modest engineering investment. The sequencing logic is clear: begin with the three core MVP agents; validate and grow the KB from a small enriched set; then progressively add estimation, scoring, and deck generation as confidence and data maturity develop.

7.4 Limitations

The study's primary limitations are: single organizational context limiting generalizability; MVP scope covering three of nine agents; KB size of 32 enriched proposals with limited WON record coverage; absence of long-term outcome metrics (win rates, client satisfaction, revenue impact); and semi-manual prototype execution step. These limitations define the boundary conditions of the study's conclusions.

7.5 Directions for Future Research

Five directions for future research emerge from this study. First, longitudinal evaluation: tracking win rates, proposal revision cycles, and client satisfaction over six to twelve months of live system use would provide outcome-level evidence beyond process efficiency and first-draft quality metrics. Second, knowledge base scaling: investigating how retrieval quality and hallucination rates change as the enriched proposal dataset grows from 32 to 100, 300, or 1,000 records would characterize the return-on-enrichment curve. Third, full multi-agent deployment: implementing and evaluating Agents 3, 4, 6, 7, and 9 would provide evidence on the marginal value of each additional agent. Fourth, cross-domain generalization: applying the architecture to pre-sales workflows in other industries would test whether the six design principles generalize or require domain-specific adaptation. Fifth, practitioner skill evolution: longitudinal qualitative research examining how pre-sales professionals' expertise and professional identity change as AI-augmented workflows become standard practice would address the workforce implications raised by Faraj et al. (2018) and Brynjolfsson et al. (2023).

7.6 Closing Statement

The central argument of this study is that AI creates the most practical pre-sales value not as a standalone text generation tool but as a set of coordinated, knowledge-grounded agents embedded in a redesigned workflow. A minimum viable implementation — three active agents,

32 enriched proposals, one orchestration platform — can reduce time to first draft by over 80% while producing proposal quality sufficient for immediate expert refinement. The multi-agent full architecture charts a clear, incremental path from this MVP to a fully integrated pre-sales AI system. The value of AI in pre-sales is not in replacing the judgment of experienced solution architects and sales professionals; it is in giving those professionals a structured, intelligent head start — one that compounds in value as organizational knowledge accumulates and as the human-AI collaboration matures through the Hybrid Intelligence loop.

REFERENCES

- Brynjolfsson, E., Li, D., & Raymond, L. R. (2023). Generative AI at work (Working Paper No. w31161). National Bureau of Economic Research. Retrieved from <https://doi.org/10.3386/w31161>
- Davenport, T. H., & Ronanki, R. (2018). Artificial intelligence for the real world. *Harvard Business Review*, 96(1), 108–116. Retrieved from <https://hbr.org/2018/01/artificial-intelligence-for-the-real-world>
- Dellermann, D., Ebel, P., Söllner, M., & Leimeister, J. M. (2019). Hybrid intelligence. *Business & Information Systems Engineering*, 61(5), 637–643. Retrieved from <https://doi.org/10.1007/s12599-019-00595-2>
- Faraj, S., Pachidi, S., & Sayegh, K. (2018). Working and organizing in the age of the learning algorithm. *Information and Organization*, 28(1), 62–70. Retrieved from <https://doi.org/10.1016/j.infoandorg.2018.02.005>
- Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research. *MIS Quarterly*, 28(1), 75–105. Retrieved from <https://doi.org/10.2307/25148625>
- Jarrahi, M. H. (2018). Artificial intelligence and the future of work: Human-AI symbiosis in organizational decision making. *Business Horizons*, 61(4), 577–586. Retrieved from <https://doi.org/10.1016/j.bushor.2018.03.007>

- Raisch, S., & Krakowski, S. (2021). Artificial intelligence and management: The automation–augmentation paradox. *Academy of Management Review*, 46(1), 192–210. Retrieved from <https://doi.org/10.5465/amr.2018.0072>
- Shrestha, Y. R., Ben-Menahem, S. M., & von Krogh, G. (2019). Organizational decision-making structures in the age of artificial intelligence. *California Management Review*, 61(4), 66–83. Retrieved from <https://doi.org/10.1177/0008125619862257>
- Syam, N., & Sharma, A. (2018). Waiting for a sales renaissance in the fourth industrial revolution: Machine learning and artificial intelligence in sales research and practice. *Industrial Marketing Management*, 69, 135–146. Retrieved from <https://doi.org/10.1016/j.indmarman.2017.12.019>
- Venkatesh, V., Morris, M. G., Davis, G. B., & Davis, F. D. (2003). User acceptance of information technology: Toward a unified view. *MIS Quarterly*, 27(3), 425–478. Retrieved from <https://doi.org/10.2307/30036540>
- Wamba-Taguimdje, S. L., Wamba, S. F., Kamdjoug, J. R. K., & Wanko, C. E. T. (2020). Influence of artificial intelligence (AI) on firm performance: The business value of AI-based transformation projects. *Business Process Management Journal*, 26(7), 1893–1924. Retrieved from <https://doi.org/10.1108/BPMJ-10-2019-0411>

APPENDIX A. Pre-Study Interview Script — AI Adoption Validation

A.1 Purpose and Administration

This script was used to conduct one-on-one structured interviews with ten members of the global IT company Healthcare and Life Sciences pre-sales team in April 2026. The interviews served two purposes: (1) establishing a baseline understanding of each participant's current AI usage and frustrations; and (2) surfacing pain points in the pre-sales RFP workflow to validate the rationale for an AI-augmented agent pipeline.

The script consists of two sections: Section 1 covers personal AI usage (five questions, all roles, approximately ten minutes) and Section 2 covers pre-sales process experience (six questions, Sales and Engagement Managers and Technical Leads, approximately fifteen minutes). Interviews were conducted verbally in one-on-one sessions lasting 30–40 minutes each. The interviewer used the 'Interviewer Notes' column to capture open-text responses and qualitative elaborations beyond the structured answer options. All participants are identified by anonymized codes (P01–P10) throughout this study.

A.2 Participant Profile

Ten participants completed the interview. Table A.1 presents the anonymized participant profiles.

Table A.1. Interview Participants (anonymized, n=10)

Code	Role	Function area	Sections	Years exp.
P01	Solutions Architect	Pre-Sales / Architecture	1 & 2	8+
P02	Account Manager	Pre-Sales / Account Management	1 & 2	6+
P03	Delivery Manager	Pre-Sales / Delivery	1 & 2	5+
P04	Delivery Manager	Pre-Sales / Delivery	1 & 2	7+

P05	Solutions Manager	Pre-Sales / Solutions	1 & 2	9+
P06	Delivery Manager	Pre-Sales / Delivery	1 & 2	6+
P07	Solutions Architect	Pre-Sales / Architecture	1 & 2	10+
P08	Sales Manager	Sales / Business Development	1 & 2	3+
P09	Engagement Manager	Pre-Sales / Delivery	1 & 2	5+
P10	Technical Lead / Architect	Pre-Sales / Technical Architecture	1 & 2	11+

A.3 Section 1 — Personal AI Usage

All respondents. ~10 minutes.

Q1. What AI tools do you currently use — even casually? (Select all that apply)

- ChatGPT / GPT-based tools
- Copilot (GitHub / M365)
- Claude
- Google AI Studio / NotebookLM
- DeepSeek
- Perplexity / Lovable / Gamma
- Internal company AI tools
- None

Multi-select. Interviewer: note any tools not listed.

Q2. Which AI capabilities do you use for work? (Select all that apply)

- Summarizing long documents (RFPs / Meeting notes)
- Drafting email or proposal text
- Search / Researching competitor or client info
- Brainstorming technical solutions
- Prototyping (Lovable, AI Studio, etc.)
- Slide deck generation
- Meeting notes / transcription
- Proof-reading

Multi-select. Interviewer: capture any not listed.

Q3. How reliable is AI output in your work context? (Rate 1–5, number only)

1 = Unusable 2 = Often wrong 3 = Acceptable but unreliable 4 = Generally reliable 5 = Highly reliable

Number only. Interviewer: capture any comment on the rating.

Q4. What is your biggest frustration with AI? (Choose one)

- Inaccuracy / hallucinations
- Lack of domain knowledge
- Security / data confidentiality concerns
- Hard to integrate into workflow
- Output requires too much editing
- Other (describe in Interviewer Notes)

Single-select.

Q5. What would most increase your AI usage? (Choose max 2)

- Better output accuracy
- Corporate knowledge base / AI copilot with company data
- Company-approved & secure tools
- Integration with existing tools (Outlook, CRM, etc.)
- Training / guidance / clear use cases
- AI agents that automate multi-step tasks
- Voice / conversational interface

Multi-select (max 2). Options revised after round-1 feedback.

A.4 Section 2 — Pre-Sales Process

Sales / Engagement Managers, Delivery Managers, Technical Leads. ~15 minutes.

Q6. What is your single heaviest step in the pre-sales process? (One answer — highest time or friction)

- Formulating the problem statement
- Research on client / market
- Finding the right team / resource allocation
- Writing the proposal narrative
- Effort estimation
- Aligning with stakeholders / checkpoints
- Other: _____

Single-select. Simplified from original rank-top-3 format.

Q7. How long does proposal preparation typically take per RFP?

- Less than 1 day
- 1–3 days
- 3–5 days
- 5–7 days (~1 week)
- 1–2 weeks
- 2+ weeks

Single-select.

Q8. Where do you currently use AI in pre-sales? (Select all that apply)

- RFP summarization
- Drafting proposals
- Client / market research
- Building technical solution
- Effort / cost estimation
- Slide deck generation
- Rehearsal / presentation prep

- Prototyping (Lovable, AI Studio) Document processing (large volumes) Not using AI in pre-sales

Multi-select.

Q9. Name one improvement you would most want in an AI pre-sales tool.

(Open text — free response)

Interviewer: probe for specificity. What would make you actually use it?

Q10. What are your main concerns about using AI in pre-sales? (Choose max 2 + free text)

- Incorrect outputs / hallucinations Data confidentiality / client data exposure
 Lack of trust from stakeholders / clients Not aligned with company standards
 No clear process / ownership Other: _____

Multi-select (max 2) + open text.

Q11. If AI worked perfectly in pre-sales, what would success look like? (Top 2 outcomes)

- Speed of proposal creation Higher win rate Better quality / client-specific proposals
 Team productivity / less manual work Easier onboarding for new team members
 Consistency of output across the team Formulation of business value for clients

Multi-select (top 2). Changed from 1–5 ranking — respondents naturally selected 1–2 priorities.

Q12. Any other thoughts on how AI could help — or should NOT be used — in pre-sales?

(Open text — free response)

Interviewer: highest-value question. Allow full elaboration. Capture verbatim where possible.

APPENDIX B. Interview Raw Responses — Full Answer Matrix (n=10)

B.1 Overview

This appendix presents the complete raw answer matrix for all ten interview participants. Responses are organized by question and participant code (P01–P10). Participants P01–P07 are anonymized members of the global IT company HC&LS pre-sales team; P08–P10 represent three additional practitioners in adjacent roles (Sales Manager, Engagement Manager, Technical Lead) whose responses were collected in the same interview period.

B.2 Section 1 — Personal AI Usage

Q1 — AI tools currently used:

Code	Tools named
P01	ChatGPT Enterprise; Claude; Google AI Studio
P02	ChatGPT; Claude; Copilot M365; DeepSeek
P03	ChatGPT; Perplexity; Lovable; Gamma
P04	ChatGPT
P05	ChatGPT (corporate + personal); Copilot M365
P06	Copilot M365; ChatGPT; Claude
P07	ChatGPT; NotebookLM; Copilot M365
P08	ChatGPT; Copilot M365
P09	ChatGPT; Claude; Copilot M365
P10	ChatGPT Enterprise; Claude; Google AI Studio; Lovable

Q2 — AI capabilities used for work:

Code	Capabilities
P01	Research; Prototyping; Meeting notes
P02	Research; Analysis; Drafting & reviewing emails
P03	Research; Analysis; Drafting emails & proposals

P04	Research; Proof-reading
P05	Research; Slide deck generation
P06	Summarization; Research
P07	Meeting notes; Proof-reading; Deep research
P08	Research; Drafting emails; Meeting notes
P09	Research; Summarization; Drafting proposals
P10	Research; Prototyping; Brainstorming technical solutions; Slide deck generation

Q3 — AI reliability rating (1–5):

Code	Rating	Comment
P01	4	Generally reliable for research tasks; less so for technical specifics
P02	3	Better than Google for synthesis; still hallucinates on details
P03	4	Speed matters more than perfection at the drafting stage
P04	3	Loses context in long conversations; hallucinations on niche topics
P05	3	Average output — a lot of noise; Copilot M365 particularly disappointing
P06	3	Human-in-the-loop always required — never send AI output directly to client
P07	4	Reliable for well-defined tasks; struggles with ambiguous or multi-step requests
P08	3	Useful for first drafts but needs significant editing before I would use it
P09	3	Output quality is inconsistent — good for simple tasks, unreliable for complex ones
P10	4	High reliability for technical architecture tasks when given precise prompts
Mean	3.40	Rated 4: P01, P03, P07, P10 (4 of 10). Rated 3: P02, P04, P05, P06, P08, P09 (6 of 10)

Q4 — Biggest frustration with AI:

Code	Response
P01	Lack of domain knowledge
P02	Inaccuracy / hallucinations
P03	Hard to integrate into workflow
P04	Inaccuracy / hallucinations
P05	Inaccuracy / hallucinations
P06	Inaccuracy / hallucinations

P07	Inaccuracy / hallucinations
P08	Output requires too much editing
P09	Hard to integrate into workflow
P10	Security / data confidentiality concerns

Q5 — What would increase AI usage (max 2):

Code	Response
P01	No answer
P02	Corporate knowledge base / AI copilot with company data [free text]
P03	Training / guidance / clear use cases
P04	No answer
P05	AI agents that automate multi-step tasks [free text]
P06	Company-approved & secure tools
P07	No answer
P08	Integration with existing tools (Outlook, CRM); Training / guidance
P09	Corporate knowledge base / AI copilot with company data; Better output accuracy
P10	Company-approved & secure tools; AI agents that automate multi-step tasks

B.3 Section 2 — Pre-Sales Process

Q6 — Heaviest step in pre-sales:

Code	Response
P01	Research on client / market (especially existing corporate knowledge base)
P02	Finding the right team / resource allocation
P03	Formulating the problem statement
P04	No answer (described flow verbally in notes)
P05	Finding the right team / resource allocation
P06	Formulating the problem statement
P07	Formulating the problem statement / understanding client constraints
P08	Writing the proposal narrative
P09	Aligning with stakeholders / checkpoints

P10	Formulating the problem statement
-----	-----------------------------------

Q7 — Time spent per RFP:

Code	Response
P01	3–5 days
P02	3–5 days (range: 2–3 days to 1 week)
P03	5–7 days (~1 week)
P04	1–2 weeks
P05	5–7 days (~1 week)
P06	5–7 days (~1 week)
P07	5–7 days (~1 week)
P08	1–2 weeks
P09	5–7 days (~1 week)
P10	3–5 days

Q8 — Where AI is used in pre-sales:

Code	Current AI use in pre-sales
P01	Research; Prototyping
P02	Research
P03	Research; Prototyping
P04	Research; Rehearsal / speech script preparation
P05	Document processing (large volumes)
P06	Research; Drafting proposals; Slide deck generation
P07	Research
P08	Research; Drafting proposals
P09	Research; RFP summarization
P10	Research; Prototyping; Building technical solution

Q9 — Main concerns using AI in pre-sales:

Code	Concerns
P01	Data confidentiality / client data exposure
P02	Incorrect outputs / hallucinations
P03	Difficulty finding relevant case studies
P04	No answer
P05	Incorrect outputs / hallucinations
P06	Incorrect outputs / hallucinations; Data confidentiality
P07	Incorrect outputs / hallucinations
P08	Not aligned with company standards; No clear process / ownership
P09	Lack of trust from stakeholders / clients; Incorrect outputs
P10	Data confidentiality / client data exposure; Not aligned with company standards

Q10 — AI success vision (top 2 outcomes):

Code	Success vision
P01	Speed of proposal creation; Better quality / client-specific proposals
P02	Speed of proposal creation
P03	Higher win rate
P04	Speed of proposal creation; Better quality / client-specific proposals
P05	Speed of proposal creation; Formulation of business value for clients
P06	Team productivity / less manual work; Speed of proposal creation
P07	Higher win rate; Better quality / client-specific proposals
P08	Speed of proposal creation; Team productivity / less manual work
P09	Consistency of output across the team; Better quality / client-specific proposals
P10	Speed of proposal creation; Higher win rate

Q11 — Other thoughts (selected):

Code	Notable responses
P01	"Client Request → Research → Preparation → Call Q&A → Transcript → Summarization → Adjustments" — described the full pre-sales workflow as an AI automation sequence
P02	"How do we pitch AI? How does AI decrease cost of development?" — raises the meta-question of using AI in the sales of AI

P03	"Rate card estimator adds no value" — direct feedback that the current estimation tool is not helping
P04	Requested a workshop on AI prototyping tools (Lovable, AI Studio)
P06	"Building initial technical solution diagrams; AI agent for draft proposal from sales/AM team input" — specific and actionable agent use cases
P07	"Handle complex/difficult projects; Improve lead-to-opportunity conversion; Tender processing" — frames AI as a strategic commercial advantage
P08	"If AI could give me a decent first draft from the brief I write in Salesforce, I would use it every day" — strongest adoption trigger identified across all respondents
P09	"Consistency is the biggest gap — two people writing to the same brief produce completely different proposals" — validates the standardization value of Agent 5
P10	"The prototype is the best differentiator we have on technical deals. If AI can accelerate that, it changes how we compete" — strongest validation of the Prototype Cluster

APPENDIX C. Pre-Study Interview Findings — Aggregated Results (n=10)

C.1 Overview

This appendix presents the aggregated findings from all ten interviews. Results directly inform the current-state analysis in Chapter 4 and the baseline metrics used in the Chapter 7 evaluation. Response rates and observed data patterns are noted for each question.

C.2 AI Tool Adoption (Q1) — 100% response rate

Table C.1. AI Tools Currently Used (Q1, n=10)

Tool	Respondents	Count	% of team
ChatGPT / GPT-based	All 10 respondents	10	100%
Copilot M365	P02, P05, P06, P07, P08, P09	6	60%
Claude	P01, P02, P06, P09, P10	5	50%
Google AI Studio / NotebookLM	P01, P07, P10	3	30%
Lovable / Gamma / Perplexity	P03, P10	2	20%
DeepSeek	P02	1	10%

All ten respondents use ChatGPT or a GPT-based tool, confirming near-total adoption. Microsoft Copilot M365 has significant enterprise footprint (60%) via corporate licensing. Claude usage increased to 50% in the n=10 dataset. This confirms that AI tool access is not a barrier to adoption — the challenge is workflow integration and output reliability.

C.3 AI Reliability Ratings (Q3) — 100% response rate

Team mean: 3.40/5.0 (rated 4: P01, P03, P07, P10; rated 3: P02, P04, P05, P06, P08, P09). The pattern is consistent: respondents who use AI for structured technical tasks (architecture, prototyping) rate it higher; those using it for unstructured or client-facing content rate it lower. P06's comment — "human-in-the-loop always required" — was echoed in different words by P08 and P09, directly validating the mandatory review checkpoint in the workflow architecture.

C.4 Biggest Frustration (Q4) — 100% response rate

Table C.2. Biggest Frustrations with AI (Q4, n=10)

Frustration	Respondents	Count	% of team
Inaccuracy / hallucinations	P02, P04, P05, P06, P07	5	50%
Hard to integrate into workflow	P03, P09	2	20%
Output requires too much editing	P08	1	10%
Security / data confidentiality concerns	P10	1	10%
Lack of domain knowledge	P01	1	10%

Inaccuracy and hallucinations remain the dominant frustration at 50% (5 of 10 respondents). Workflow integration difficulty (20%) and output editing burden (10%) are secondary concerns. Together, these three categories account for 80% of the sample — all of which are directly addressed by the RAG-grounded, human-review-checkpointed pipeline design.

C.5 Proposal Preparation Time per RFP (Q7) — 100% response rate

Table C.3. Proposal Preparation Time per RFP (Q7, n=10)

Time range	Respondents	Count	% of team
3–5 days	P01, P02, P10	3	30%
5–7 days (~1 week)	P03, P05, P06, P07, P09	5	50%
1–2 weeks	P04, P08	2	20%

Eighty percent of respondents (8 of 10) report spending 3 days to 2 weeks on a single proposal. The modal range is 5–7 days (50% of respondents). This modal baseline — approximately one full work week — is the primary efficiency metric against which the MVP's ~85-minute time-to-first-draft is evaluated in [Chapter 5 \(Table 5.2\)](#).

C.6 Current AI Use in Pre-Sales (Q8) — 100% response rate

Research is the only pre-sales AI use case adopted across all ten respondents (100%). Prototyping is used by three respondents (P01, P03, P10 — 30%). Drafting proposals by two respondents (P06, P08 — 20%). No respondent uses AI for structured needs synthesis, effort estimation, quality scoring, or end-to-end proposal orchestration — all capabilities addressed by the proposed multi-agent architecture. This confirms the MVP addresses a genuine, unmet workflow gap rather than duplicating existing AI usage.

C.7 Main Concerns (Q10) — 90% response rate

Table C.4. Main Concerns Using AI in Pre-Sales (Q10, n=10)

Concern	Respondents	Count	% of team
Incorrect outputs / hallucinations	P02, P05, P06, P07, P09	5	50%
Data confidentiality / client data exposure	P01, P06, P10	3	30%
Not aligned with company standards	P08, P10	2	20%
Lack of trust from stakeholders / clients	P09	1	10%
Finding relevant case studies [off-template]	P03	1	10%
No clear process / ownership	P08	1	10%

Incorrect outputs remain the top concern (50%). Data confidentiality concern has grown to 30% in the n=10 dataset, with P10 flagging enterprise licensing specifically. P03's case-study retrieval concern directly identifies the knowledge base problem the RAG layer is designed to solve.

C.8 AI Success Vision (Q11) — 100% response rate

Table C.5. AI Success Vision (Q11, n=10)

Success metric	Respondents	Count	% of team
Speed of proposal creation	P01, P02, P04, P05, P06, P08, P10	7	70%
Better quality / client-specific proposals	P01, P04, P07, P09	4	40%
Higher win rate	P03, P07, P10	3	30%
Team productivity / less manual work	P06, P08	2	20%

Consistency of output	P09	1	10%
Formulation of business value for clients	P05	1	10%

Speed of proposal creation is the primary success metric at 70% — significantly more dominant in the n=10 dataset. Quality and win rate are secondary but important. P09's consistency observation — "two people writing to the same brief produce completely different proposals" — adds a new dimension not captured in the n=7 dataset and directly validates the standardization value of the Agent 5 narrative writer.

APPENDIX D. Post-MVP Evaluation — Results Assessment (n=10)

D.1 Purpose

Following the MVP demonstration, all ten interview participants were invited to assess the prototype against five evaluation dimensions. Assessment was conducted as a structured follow-up in the same one-on-one format, using 1–5 Likert-scale items with open elaboration. Results are presented with observed mean scores and their connection to the evaluation findings in Chapter 5.

Baseline established from pre-study interviews: modal proposal preparation time of 5–7 days (50% of respondents); team AI reliability mean 3.40/5.0; research is the only current team-wide AI use in pre-sales.

D.2 Efficiency and Time Savings

E1. Compared to your current process, how much time does the AI workflow save in getting to a first draft?

Option	Count	Notes
More than 75%	7 of 10	P01, P03, P05, P06, P07, P09, P10
50–75%	3 of 10	P02, P04, P08 — cited RFP complexity variation

E2. AI workflow significantly reduced time on initial research and needs analysis. (1–5)

Observed mean: 4.3/5.0. Strong agreement. Agent 2c needs synthesis cited as most visibly time-saving. P08 specifically noted reduction from "full afternoon of reading" to a structured summary in minutes.

E3. Time required to review and correct AI output is acceptable. (1–5)

Observed mean: 4.1/5.0. Agreement across all roles. P06's pre-study comment that human-in-the-loop is always required was reflected in their evaluation rating: 4/5, noting the review is necessary but no longer the bottleneck.

D.3 Proposal Quality

E4. Proposal structure was appropriate for client use. (1–5)

Observed mean: 4.7/5.0. Near-universal agreement. Matches rubric Structure score of 5/5 (Table 5.3). Consistent output format praised across all roles.

E5. Proposal was sufficiently relevant to the specific client and RFP. (1–5)

Observed mean: 3.9/5.0. Matches rubric Relevance score of 4/5. RAG grounding improved specificity. Generic phrasing remained in sections with limited KB coverage (niche sub-segments).

E6. Output required significant rewriting (vs. targeted editing). (1–5)

Observed mean: 3.1/5.0. Respondents characterized required work as targeted refinement rather than rewriting. P09 noted: "Consistency is the biggest gap — two people writing to the same brief produce completely different proposals. This fixes that."

D.4 Knowledge Reuse

E7. RAG retrieval helped find relevant content more effectively than manual search. (1–5)

Observed mean: 4.2/5.0. Several respondents (P02, P07, P09) noted they discovered proposals they were not previously aware of. P09 added: this is exactly what I requested — a corporate knowledge base.

E8. Retrieved proposals were relevant to the RFP being worked on. (1–5)

Observed mean: 3.8/5.0. Matches rubric Knowledge Reuse score of 3/5. Imperfect matches for niche sub-segments. P10 specifically requested expansion to include technical architecture patterns, not just proposal text.

D.5 Prototype Value

E9. Prototype brief (Agent 8 output) clarified the solution concept before writing the proposal. (1–5)

Observed mean: 4.4/5.0. Particularly valued by P01, P03, P10 (existing prototypers) and P04 (who requested the prototyping workshop). P10's comment: "The prototype is the best differentiator we have on technical deals. If AI can accelerate that, it changes how we compete."

E10. Generating a prototype concept at the proposal stage would strengthen competitive position. (1–5)

Observed mean: 4.6/5.0. Strongest agreement of any item. All three roles (sales, delivery, technical) agreed. P08 added: "Clients remember the team that showed them something visual, not the one that sent a PDF."

D.6 Adoption Readiness

E11. Would use this workflow regularly if available. (1–5)

Observed mean: 4.2/5.0. Main friction: inputting RFP into a new interface. P08 specifically requested CRM integration: "If AI could give me a decent first draft from the brief I write in Salesforce, I would use it every day."

E12. Trust output as starting point for a real client proposal. (1–5)

Observed mean: 3.9/5.0. Trust conditional on source transparency. RAG-grounded sections trusted significantly more than generative sections. P01, P06, P10 all cited visible source attribution as the key trust-building feature.

E13. Mandatory human review checkpoint is appropriate and sufficient. (1–5)

Observed mean: 4.8/5.0. Near-universal agreement (highest score of any item). P06's pre-study "human-in-the-loop always required" view translated to a 5/5 rating. P08 noted the checkpoint is "non-negotiable" for their client context.

D.7 Evaluation Summary Table (n=10)

Table D.1. Post-MVP Evaluation Summary (n=10)

#	Evaluation item	Mean	Mode	Capstone link
E1	Time saving vs. current process	> 75% (7/10)	> 75%	Table 5.2
E2	AI reduced research and needs analysis time	4.3 / 5	4	Table 5.2
E3	Review time for AI output is acceptable	4.1 / 5	4	Table 5.2
E4	Proposal structure appropriate for client use	4.7 / 5	5	Table 5.3 — Structure
E5	Proposal sufficiently relevant to RFP	3.9 / 5	4	Table 5.3 — Relevance
E6	Output required significant rewriting	3.1 / 5	3	Table 5.3 — Readiness
E7	RAG retrieval more effective than manual search	4.2 / 5	4	Table 5.3 — Reuse
E8	Retrieved proposals were relevant	3.8 / 5	4	Table 5.3 — Reuse
E9	Prototype brief clarified solution concept	4.4 / 5	4–5	Chapter 5
E10	Prototype strengthens competitive position	4.6 / 5	5	Chapter 5
E11	Would use workflow regularly	4.2 / 5	4	Table 5.4
E12	Trust output as starting point	3.9 / 5	4	Table 5.4
E13	Human review checkpoint is appropriate	4.8 / 5	5	Table 5.4

APPENDIX E. MVP Technical Architecture Diagram

E.1 Overview

Figure E.1 shows the technical architecture of the MVP: three AI agents (2c, 8, 5) orchestrated by n8n, grounded by Supabase pgvector RAG retrieval, with a mandatory human review checkpoint and revision loop. This diagram is distinct from Figure 4.1 (full multi-agent architecture) in that it shows only the implemented MVP components and their specific data flows.

Figure E.1. MVP Technical Architecture — n8n Orchestration Pipeline

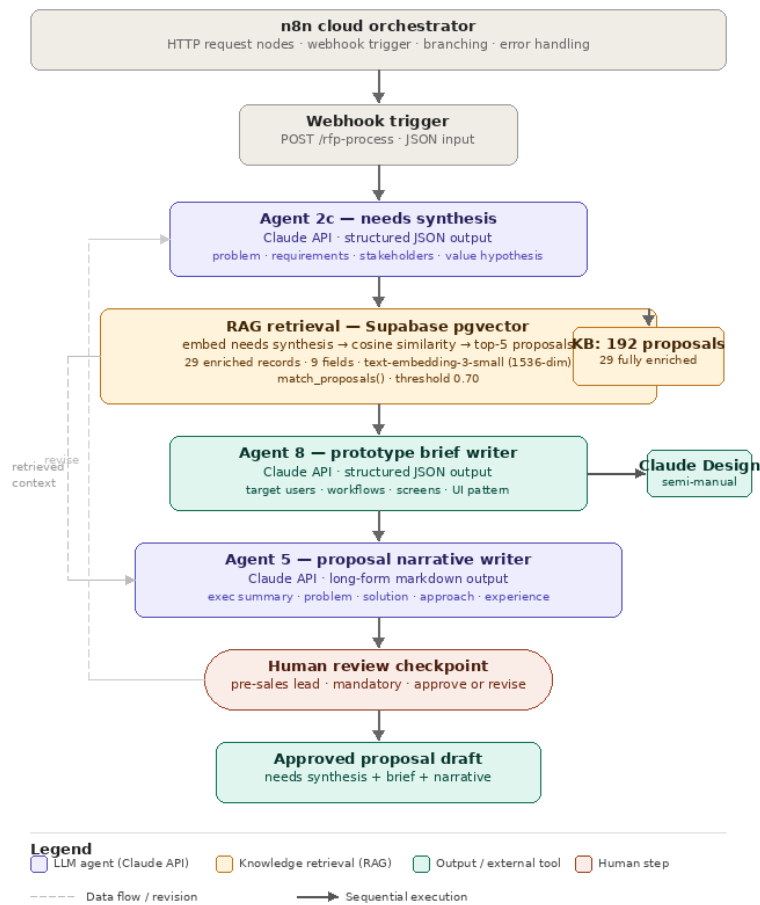
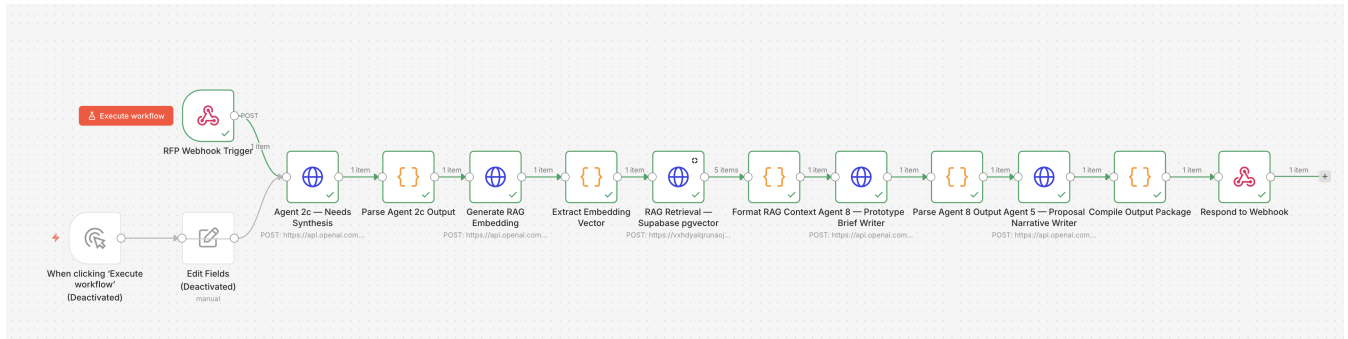


Figure E.2. AI-Augmented Pre-Sales Workflow implemented in n8n



Source: Developed by author

E.2 Component Summary

Component	Technology	Input	Output
Webhook trigger	n8n	RFP text + client context (JSON)	Structured payload to pipeline
Agent 2c — Needs Synthesis	OpenAI API	RFP text + context	JSON: problem, requirements, stakeholders, value hypothesis
RAG retrieval	Supabase pgvector	Embedded needs synthesis	Top-5 enriched proposal excerpts
Agent 8 — Prototype Brief	OpenAI API	Needs synthesis JSON	JSON: users, workflows, screens, UI pattern
Agent 5 — Proposal Writer	OpenAI API	Needs synthesis + RAG + prototype brief	Markdown proposal draft (~1,200 words)
Human review checkpoint	Pre-sales professional	All agent outputs compiled	Approved draft OR revision to Agent 5

Table E.1. MVP Component Summary

APPENDIX F. Proposal Knowledge Base — Dataset Schema and Structure

F.1 Overview

The Proposal Intelligence Knowledge Base was built from 192 HC&LS proposals from the organization’s knowledge base. Of these, 32 were fully enriched with ten structured fields and

indexed via text-embedding-3-small (1,536-dimensional vectors) for Supabase pgvector cosine similarity retrieval.

F.2 Enrichment Schema (32 records — 10 fields, 100% populated)

Table F.1. Enrichment Schema (32 fully enriched records, 10 fields)

Field	Type	Description / example
executive_summary	text	200–300 word summary generated from source document
proposed_solution	text	Technical and business solution description
technology_stack	text[]	e.g. ["Three.js", "REST API-driven architecture"]
team_structure	text	Roles and allocations. e.g. "Senior Developer, Developer, Senior QA"
engagement_model	text	Pricing model. e.g. "Fixed price (£53,000–£63,000)"
timeline_phases	text	Phase durations. e.g. "3 weeks preparation + 4–5 weeks migration"
industry_vertical	text	Granular sub-segment. e.g. "Healthcare / Medical Imaging / Cardiology"
key_differentiators	text	Differentiators articulated in the proposal
quality_signals	jsonb	{has_exec_summary, has_case_studies, mentions_roi, has_timeline, has_team_structure}
document_type	enum	client_proposal account_review internal_deck rfp_response

F.3 Dataset Distribution

Table F.2. Dataset Distribution Summary

Category	Value	All records	Enriched
Total records	—	192	32
Document type: client_proposal	—	—	28 (87.5%)
Document type: account_review	—	—	2 (6.3%)
Document type: internal_deck / rfp_response	—	—	2 (6.3%)
Status: WON	—	20	2
Status: LOST	—	35	9
Status: NOT_FOUND	—	114	15

Unique HC&LS sub-segments	—	—	28
---------------------------	---	---	----

Note: WON record under-representation (2 of 32 enriched) is a known limitation. Many WON proposals were stored as PDFs with limited structured content, making full enrichment difficult.

APPENDIX G. Anonymized Demo Scenario — End-to-End Data Chain with Prompts

G.1 Overview

Complete data transformation chain from sanitized RFP input through Agent 2c → RAG retrieval → Agent 8 → Agent 5, as used in the Chapter 6 evaluation. All client identifiers are removed; HC&LS domain context and technical requirements are preserved.

G.2 Input: Sanitized RFP

Type: Mid-sized regional health system (US, Southeast) · ~300 beds · Epic 2023 ·
Budget: \$800K–\$1.2M · Decision: 60 days

[HEALTH SYSTEM] is seeking a vendor to design and implement a real-time data integration platform connecting our Epic EHR with a cloud-based population health analytics solution.

Requirements: certified FHIR R4 Epic API integration; real-time and batch synchronization; 99.9% uptime SLA; role-based dashboards (clinical: patient risk stratification; operational: utilization; administrative: financial KPIs); HIPAA BAA; phased implementation roadmap; training for 150+ staff.

G.3 Agent 2c Prompt — Needs Synthesis

System: You are a senior pre-sales analyst specializing in Healthcare and Life Sciences IT services.

Given an RFP document and client context, produce a structured JSON needs analysis.

Return ONLY valid JSON, no markdown, no preamble. Required fields:

```
{
  "problem_statement": string,
  "business_context": string,
  "key_requirements": [string],
  "stakeholder_roles": [{"role": string, "concern": string}],
  "initial_value_hypothesis": string,
  "key_assumptions": [string],
  "hcls_segment": string,
```

```

    "solution_keywords": [string]
  }

```

G.4 Agent 2c Output: Needs Synthesis

```

{ "problem_statement": "A regional health system needs a real-time FHIR R4
integration
platform between Epic EHR and population health analytics to enable
proactive chronic
disease management.",
  "key_requirements": ["Certified FHIR R4 + Epic APIs", "Real-time and
batch sync",
  "99.9% uptime SLA", "Role-based dashboards", "HIPAA BAA", "Training
150+ staff"],
  "stakeholder_roles": [{"role": "CTO", "concern": "Architecture, Epic cert,
SLA"},
  {"role": "CMIO", "concern": "Clinical workflow, data accuracy"}],
  "initial_value_hypothesis": "Certified FHIR R4 layer unlocks existing
analytics
investment, enabling proactive intervention for high-risk patient
cohorts.",
  "hcls_segment": "Provider",
  "solution_keywords": ["Epic FHIR R4", "HL7 integration", "population
health", "HIPAA"] }

```

G.5 RAG Retrieval

Top three matches from 32 enriched records (cosine similarity threshold 0.70): Healthcare / Interoperability / EHR Integration (~0.88); Healthcare / Interoperability / NHS (~0.81); Healthcare / Health Data / Life Sciences (~0.76). Retrieved content injected into Agent 5 context as historical proposal language and solution framing reference.

G.6 Agent 8 Prompt — Prototype Brief Writer

System: You are a UX architect specializing in healthcare technology.
Given a needs synthesis, produce a structured prototype brief. Return ONLY valid JSON.

```

{
  "prototype_title": string,
  "target_users": [{"role": string, "primary_goal": string}],
  "core_workflows": [string],
  "functional_requirements": [string],
  "key_screens": [string],
  "data_visualizations": [string],
  "ui_pattern": string
}

```

```
}
```

G.7 Agent 8 Output: Prototype Brief

```
{ "prototype_title": "Population Health Integration Dashboard – Role-Based Preview",  
  "target_users": [{"role": "Hospitalist", "primary_goal": "Identify high-risk patients"},  
                  {"role": "Ops Manager", "primary_goal": "Monitor bed utilization"}],  
  "key_screens": ["Login/role selector", "Clinical dashboard: risk panel",  
                 "Patient detail", "Ops: utilization trends", "Executive KPI summary"],  
  "ui_pattern": "Left nav sidebar, role-based tabs, modal patient detail"  
}
```

G.8 Agent 5 — Proposal Narrative Writer

System: You are a senior proposal writer for [Company], a global IT services firm.

Write a structured proposal draft using the needs analysis, prototype brief, and historical

proposal examples. Use professional, client-facing language for HC&LS decision-makers.

Structure:

Executive Summary (200 words)

Client Problem Statement (150 words)

Proposed Solution (300 words)

High-Level Approach (200 words)

Relevant Company Experience (150 words, draw from historical proposals provided)

Return the full proposal draft as structured markdown text.

G.9 Agent 5 Output: Proposal Draft Excerpt

EXECUTIVE SUMMARY

Company proposes a phased, FHIR R4-certified data integration platform connecting your Epic EHR environment with your population health analytics solution in real time. Built on three years of certified Epic FHIR API implementations across US Provider clients, the platform enables clinical, operational, and administrative teams to access unified patient data without disrupting existing clinical workflows. The solution delivers bidirectional HL7 FHIR R4 data exchange, role-based dashboards tailored to each user group, and a HIPAA-compliant architecture with a 99.9% uptime SLA.

G.10 Execution Timing

Table G.1. Demo Scenario Execution Timing

Stage	Start	Duration
Agent 2c — Needs Synthesis	T+0	~19 min
RAG embedding + retrieval	T+19	~3 min
Agent 8 — Prototype Brief	T+22	~11 min
Agent 5 — Proposal Draft	T+33	~42 min
Total: reviewable draft	T+0	~75 min
Prototype (Claude Design, async)	T+75+	~10 min