Volume 5, No. 11 November 2026 - (1212-1228)
p-ISSN 2980-4868 | e-ISSN 2980-4841
https://ajesh.ph/index.php/gp

Designing a Trustworthy AI-Enabled Knowledge Management Architecture to Improve Sales Readiness in B2B Technology Services

Gerry Gusman Wibowo*, Heru Purnomo Ipung, Eka Budiarto
Swiss German University, Indonesia
Email: gerry.gusman@gmail.com*

Keywords:
Knowledge Management; Design Science Research; Trustworthy AI; Retrieval-Augmented Generation; Sales Readiness

ABSTRACT

Early-stage sales readiness in business-to-business (B2B) technology services remains inconsistent due to fragmented organizational knowledge, varying levels of sales discovery capability, and the absence of standardized governance mechanisms for validating and controlling knowledge reuse. Although artificial intelligence (AI) offers opportunities to enhance knowledge retrieval, concerns regarding output reliability, hallucinations, and data integrity present significant challenges to trustworthiness. This study aimed to design a trustworthy AI-enabled knowledge management architecture to support early-stage sales readiness in B2B technology services. Adopting a Design Science Research (DSR) methodology, the proposed architecture was developed by integrating the Socialization, Externalization, Combination, and Internalization (SECI) model; the Knowledge Management Lifecycle; The Open Group Architecture Framework (TOGAF); Control Objectives for Information and Related Technologies (COBIT); the National Institute of Standards and Technology AI Risk Management Framework (NIST AI RMF); the European Union's Ethics Guidelines for Trustworthy AI; ISO/IEC 27001; and Indonesia's Personal Data Protection Law. The architecture was demonstrated through a retrieval-augmented generation (RAG) proof of concept at PT XYZ, an Indonesian information and communications technology (ICT) service provider. Data were gathered from structured and unstructured information sources, including documented organizational knowledge and tacit knowledge elicited from personnel. The proposed architecture was evaluated by six experts through focus group discussions (FGDs), using experience-weighted mean scores. The evaluation produced weighted mean scores of 4.54 for Knowledge Structure, 4.61 for Architecture, 4.61 for Governance, and 4.59 overall, all classified as Highly Achieved according to the study's evaluation criteria. These results provided expert-based support for the three proposed hypotheses concerning the architecture's assessed dimensions. The findings suggested that early-stage sales readiness depended not only on AI capabilities but also on effective knowledge management practices, structured knowledge resources, enterprise architecture alignment, governance mechanisms, information validation, and human oversight. The study concluded that organizations should prioritize centralized, well-governed knowledge repositories and adopt AI capabilities incrementally. Future research should evaluate the proposed architecture across multiple organizations and assess its long-term effects on sales readiness and business performance.

INTRODUCTION

Business-to-business (B2B) purchasing is generally more complex than business-to-consumer (B2C) purchasing because it involves organizational decision-making processes (De Keyser et al., 2025). This complexity arises from the involvement of multiple stakeholders, longer

sales cycles, and greater organizational risks, resulting in more formalized and time-consuming purchasing procedures (Webster & Wind, 1972; Wiersema, 2013). The information and communications technology (ICT) industry presents additional complexity because its products and services often require customization to meet customers' specific business requirements, IT infrastructure, and operational processes. Consequently, the value of an ICT solution depends not only on its technical features but also on its compatibility with and adaptability to the customer's organizational environment (Mittal et al., 2025).

Presales activities serve as a critical link between customers' business requirements and technical solution delivery by identifying customer needs and translating them into appropriate technical proposals (Abdelaziz & Khalil, 2023a). This function is operationally demanding because it requires the integration of technical expertise, financial considerations, and business objectives, potentially contributing to substantial workload pressures (Abdelaziz & Khalil, 2023b). Meanwhile, B2B buyers have become increasingly informed, with 66% reportedly consulting reviews before engaging with sales representatives (Gartner, 2024). This behavior increases expectations for relevant, accurate, and timely information during initial sales interactions (Andersson et al., 2024; Lundin & Kindström, 2023). These challenges are further intensified by the involvement of multiple stakeholders in organizational purchasing decisions, particularly as larger organizations typically engage more participants in software procurement processes, as summarized in Table 1.

Table 1. Number of People Involved in Software Purchase by Business Size (Gartner, 2024)

People involved	Total	Small business (5–249 employees)	Midsize business (250–2,499 employees)	Enterprise business (2,500+ employees)
Just me	12%	14%	11%	9%
2–3 people	33%	46%	21%	18%
4–6 people	32%	27%	40%	29%
7–10 people	14%	10%	16%	21%
More than 10 people	10%	4%	12%	23%

Sales and presales teams play a critical role in guiding customers through complex evaluation processes by communicating value propositions tailored to their specific business requirements. However, early-stage sales readiness often depends on individual competencies rather than structured, institutionalized knowledge management systems, resulting in inconsistent sales outcomes, uneven communication of value propositions, and inadequate opportunity qualification. Although B2B technology service organizations possess substantial knowledge assets, these resources are frequently fragmented across information systems and individual employees, limiting their accessibility and reuse (Johnston & Cortez, 2024). Artificial intelligence (AI) offers opportunities to improve knowledge retrieval and accessibility; however, concerns regarding output reliability, AI-generated hallucinations, and insufficient governance present challenges to system trustworthiness and data integrity (Alawamleh et al., 2024; Ali et al., 2024; Huang & Rust, 2021; James et al., 2025). Without clearly defined knowledge ownership, validation procedures, and governance controls, AI-enabled systems may reinforce existing inconsistencies rather than address them (Agbabiaka et al., 2025; Uren & Edwards, 2023).

A fishbone diagram, also known as an Ishikawa or cause-and-effect diagram, was used to systematically identify factors contributing to inadequate early-stage opportunity qualification. As illustrated in Figure 1, the analysis identified interrelated challenges involving people, processes,

technology, and governance. These included fragmented organizational knowledge, insufficiently standardized operational procedures, inadequate governance mechanisms, the absence of a structured capability development framework, and inconsistent customer discovery capabilities. Among these factors, knowledge fragmentation and the lack of an AI-enabled knowledge retrieval mechanism were identified as the primary structural challenges addressed in this study.

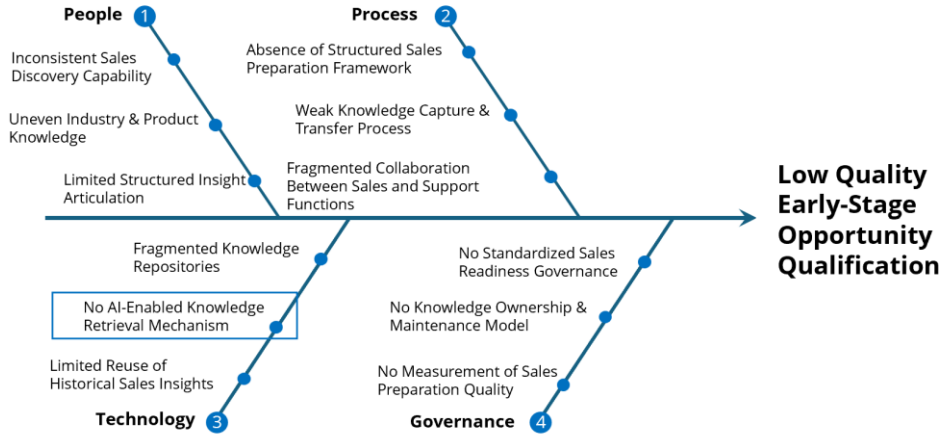


Figure 1. Fishbone Diagram of Low-Quality Early-Stage Opportunity Qualification

Accordingly, this study addressed three research problems: (1) inconsistent early-stage B2B sales readiness resulting from variations in customer discovery capabilities and insufficient structured preparation support; (2) fragmented sales knowledge assets and the absence of controlled retrieval mechanisms, limiting information accessibility and consistency; and (3) inadequate governance mechanisms for knowledge validation, ownership accountability, and sustained sales readiness. The primary objective was to design a trustworthy artificial intelligence (AI)-enabled knowledge management architecture to support early-stage sales readiness in B2B technology services. Specifically, the study aimed to identify essential knowledge domains for consistent sales preparation, design an AI-enabled architecture that integrated fragmented knowledge repositories through a unified and controlled retrieval mechanism, and develop a governance framework addressing knowledge ownership, validation procedures, access controls, and knowledge lifecycle management.

To address these objectives, three hypotheses were formulated. H1 proposed that clearly structured and accessible sales knowledge domains would improve the consistency of early-stage sales preparation. H2 proposed that a trustworthy AI-enabled knowledge retrieval architecture would improve knowledge accessibility, reliability, and controlled reuse compared with fragmented manual processes. H3 proposed that integrating structured governance mechanisms would strengthen knowledge validity, accountability, and the sustainability of sales readiness practices. The study addressed a research gap concerning the integration of AI-enabled knowledge retrieval, validation procedures, and governance mechanisms within a unified enterprise architecture for B2B technology services. Although the proposed architecture may be conceptually applicable to similar organizational settings, its evaluation was limited to a single organization and focused on architectural design and expert-based validation. The study did not include full-scale system implementation, comprehensive AI model performance testing, infrastructure optimization, or empirical measurement of downstream business outcomes, such as revenue growth or sales win rates.

METHOD

This study employed a qualitative approach using Design Science Research (DSR) as its primary methodology (Hevner et al., 2004; vom Brocke et al., 2020). The research artifact was a trustworthy AI-enabled knowledge management framework developed by integrating the Socialization, Externalization, Combination, and Internalization (SECI) model, the Knowledge Management (KM) Lifecycle, The Open Group Architecture Framework (TOGAF), Control Objectives for Information and Related Technologies (COBIT), the NIST AI Risk Management Framework (AI RMF), and the European Union's Ethics Guidelines for Trustworthy AI. The framework was demonstrated through a proof-of-concept (PoC) prototype using retrieval-augmented generation (RAG). The case organization, PT XYZ, was an Indonesian information and communications technology (ICT) service provider with more than 35 years of experience in connectivity, cloud computing, cybersecurity, collaboration, managed services, and data center solutions. Data were collected from structured sources, including customer relationship management (CRM) records, company profiles, product portfolios, internal pricing documents (Nota Dinas), and sales master data; unstructured sources, including proposals, technical documentation, and presentations; and tacit knowledge elicited from sales, presales, and technical experts. Organizational documents and focus group discussion (FGD) findings were analyzed thematically, translated into design requirements, and categorized into business, data, application, and technology architecture domains based on TOGAF. The overall research framework is presented in Figure 2.

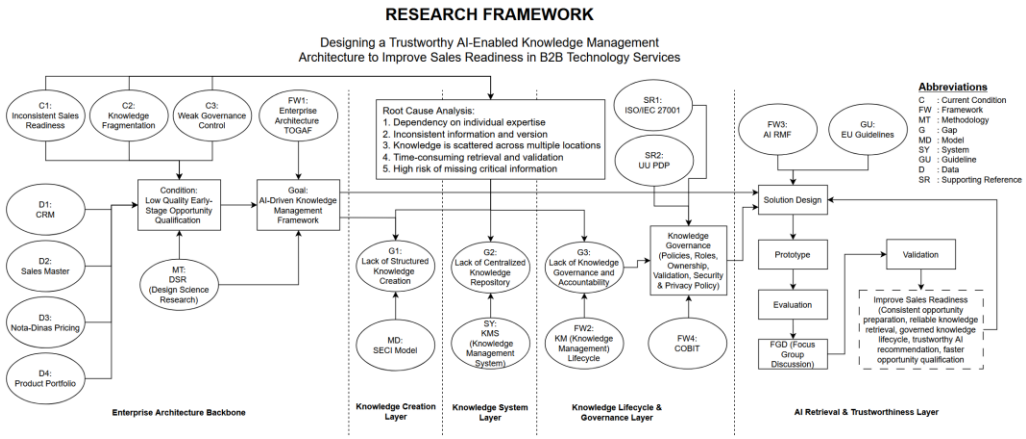


Figure 2. Integrated Research Framework for Improving Sales–Presales Readiness

The proposed framework was evaluated across three dimensions aligned with the research objectives: Knowledge Structure, informed by the SECI model and Knowledge Management Lifecycle; Architecture, based on TOGAF and retrieval-augmented generation (RAG); and Governance, guided by COBIT, the NIST AI Risk Management Framework (AI RMF), the European Union's Ethics Guidelines for Trustworthy AI, ISO/IEC 27001, and Indonesia's Personal Data Protection (PDP) Law. Expert assessments were conducted using a five-point rating scale ranging from 1 (Very Poor) to 5 (Excellent), supported by dimension-specific maturity criteria. Mean scores were classified using equal intervals of 0.80, as presented in Table 2. To account for differences in professional experience, an experience-weighted mean was calculated based on structured expert judgment principles (Cooke, 1991; Clemen & Winkler, 1999; Colson & Cooke, 2018). Validators with 10–19 years of professional experience were assigned a weight of 1.00,

while those with at least 20 years were assigned a weight of 1.25. The weighted mean was calculated as:

Table 2. Interpretation of Mean Scores

Mean Score	Interpretation
1.00–1.80	Not Achieved
1.81–2.60	Partially Achieved
2.61–3.40	Moderately Achieved
3.41–4.20	Achieved
4.21–5.00	Highly Achieved

The framework is considered successful when each of the three dimensions, and the overall mean, reaches at least 4.00 and the PoC successfully demonstrates the defined end-to-end ingestion, retrieval, and AI-assisted response workflow. Operational indicators such as retrieval time, preparation time, knowledge reuse rate, and user adoption were defined for future enterprise implementation but were not measured, as the solution was implemented as a PoC. Validation was conducted through Focus Group Discussions (FGDs) (Krueger & Casey, 2015) with experts selected purposively based on a minimum of five years of relevant experience in business, technical, or governance domains. The FGD followed five steps: introduction to the research objectives and framework, presentation of the architecture and governance model, demonstration of the PoC, individual evaluation using the structured instrument, and group discussion and recommendations. The participating experts are listed in Table 3.

Table 3. Expert Validators

Code	Expert Category	Current Position	Experience	Relevant Expertise
EV1	Enterprise Architecture Expert	Chief Executive Officer	20+ years	Enterprise Architecture, TOGAF, Business Process Integration
EV2	Business Strategy Expert	Chief Commercial Officer	20+ years	Business Strategy, Business Development, Strategic Partnerships
EV3	Presales Expert	General Manager Presales	20+ years	B2B Technology Sales, Opportunity Qualification, Sales Readiness
EV4	Knowledge Management Expert	General Manager Human Capital Management	20+ years	Knowledge Management, Organizational Learning, Capability Development
EV5	Governance Expert	Risk & Compliance Lead	10+ years	Data Governance, Risk Management, Compliance
EV6	AI Expert	AI Researcher	10+ years	Artificial Intelligence, LLM, RAG, Prompt Engineering

The PoC was limited to knowledge ingestion, the knowledge repository, vector-based semantic retrieval, AI-assisted retrieval, governance metadata and traceability, and a prototype interface, excluding production deployment, performance benchmarking, full enterprise integration, and high-availability infrastructure. It was evaluated through four business scenarios (knowledge retrieval of previous proposals, solution recommendation, tender analysis, and product information) and four trustworthy AI scenarios covering reliability and grounding, transparency and traceability, security and privacy, and human oversight and accountability.

RESULTS AND DISCUSSION

Problem Validation

The FGD findings validated three gaps. First, a significant portion of knowledge remains dependent on individual experience; project experiences, customer insights, and lessons learned are retained as personal knowledge, increasing the risk of knowledge loss when employees change roles (G1). Second, knowledge is dispersed across CRM systems, proposal archives, pricing documents, shared folders, and individual storage, making historical proposals and technical references difficult to locate (G2). Third, although ownership of certain documents has been assigned, knowledge assets are not consistently reviewed, validated, or communicated, so different stakeholders rely on different versions of the same information (G3). Together, these gaps result in inconsistent sales readiness, which is especially hard on new staff who need extra time to learn organizational capabilities and solution references. Table 4 maps each gap to its design response.

Table 4. Gap-to-Solution Mapping

Gap	Root Cause	Proposed Solution	Supporting Framework	Expected Improvement
G1 – Lack of Structured Knowledge Management	Knowledge depends on individuals and is not consistently captured or reused	Structured Knowledge Management Framework	SECI Model + Knowledge Lifecycle	Better reuse, reduced individual dependency, more consistent preparation
G2 – Lack of a Centralized Knowledge Repository	Knowledge is scattered across multiple systems and documents	Centralized repository with AI-enabled retrieval	Knowledge Repository + RAG + AI Retrieval Layer	Faster access, reduced manual search, improved reuse
G3 – Lack of Knowledge Governance and Accountability	Ownership, validation, maintenance, security, and privacy are inconsistently applied	Governance and lifecycle framework with ownership, validation, access control, traceability, and human oversight	COBIT + NIST AI RMF + EU Guidelines + ISO/IEC 27001 + PDP Law + KM Lifecycle	Improved reliability, accountability, trustworthiness, security, and sustainability

Framework Development

The framework was developed iteratively through four TOGAF-based phases: Architecture Vision, Business Architecture, Architecture Design, and Prototype Development. The Architecture Vision aligned the validated gaps with stakeholders (Sales, Presales, Product, Knowledge Owner, Governance Team, and Management) and with expected business value: faster retrieval, improved readiness, and knowledge reuse, balanced against the costs of repository maintenance and AI infrastructure and the risks of outdated content, AI inaccuracies, and adoption challenges. In the Business Architecture, the research boundary was set on the Prospecting and Value Proposition stages of the opportunity generation process, where Qualification, Solutioning, and Pricing activities are the most knowledge-intensive (Figure 3).

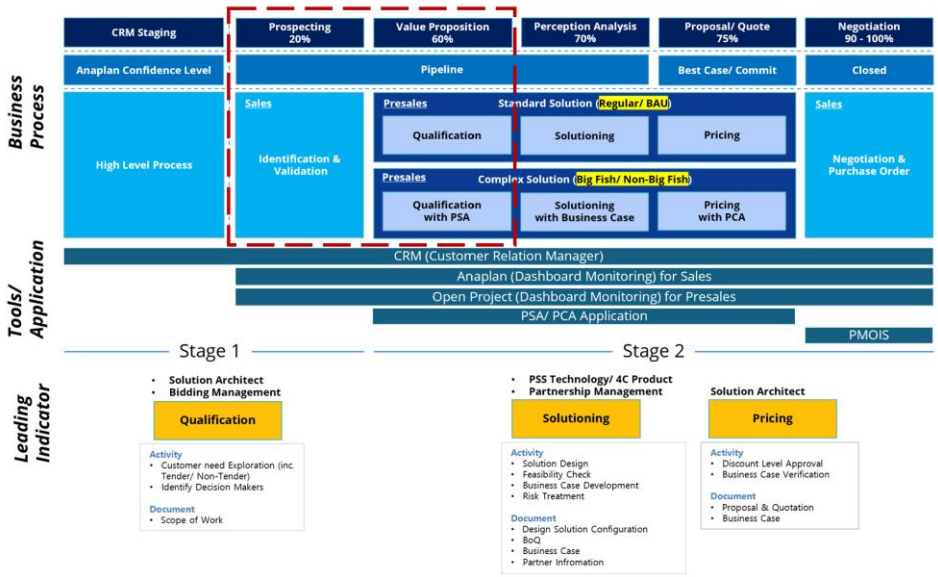


Figure 3. Opportunity Generation Process and Research Boundary

In the current model, knowledge from account planning, proposal repositories, shared folders, personal knowledge, and pricing files is fragmented, requiring manual search, collection, validation, and synthesis before value proposition activities. The proposed model (Figure 4) consolidates knowledge sources into a centralized repository, governed by a layer for ownership, validation, quality control, access control, and lifecycle management, and accessed through an AI Retrieval Layer (RAG and LLM) to support qualification, solutioning, and pricing without altering the existing business workflow. Sales can thus access basic product and solution knowledge more independently, while Presales and Product focus on complex requirement analysis and solution design. AI supports retrieval rather than replacing human expertise.

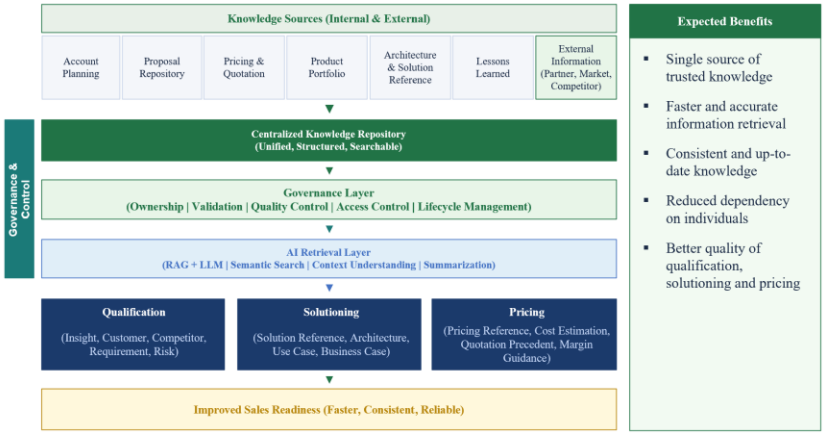


Figure 4. Proposed Knowledge Support Model (To-Be)

In the Architecture Design phase, the Data Architecture (Figure 5) transforms organizational knowledge from internal documents, operational data, employee knowledge, media, and external sources through extraction, cleansing, transformation, chunking, and embedding into a centralized repository, supported by metadata management, an information layer, and governance controls such as role-based access, audit logging, privacy management, and data lifecycle.

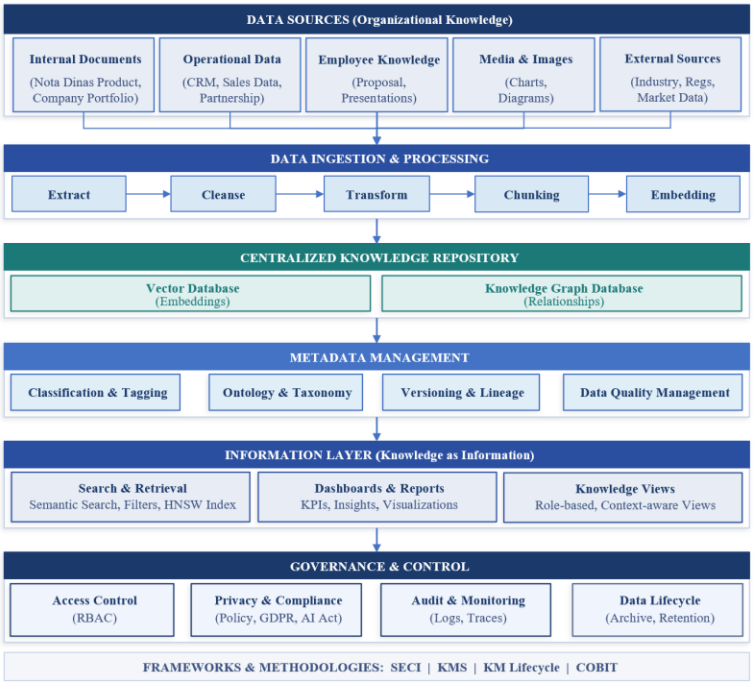


Figure 5. Data Architecture

The Application Architecture (Figure 6) consists of five logical layers: User Interface, Application Services, AI Service, Integration, and Application Management. Trustworthiness principles are embedded throughout rather than as standalone components: transparency through source citation, reliability through retrieval from validated sources, human oversight through review and approval workflows, accountability through audit logging and ownership assignment, and traceability through source tracking and retrieval history, aligning with the NIST AI RMF and EU Ethics Guidelines.

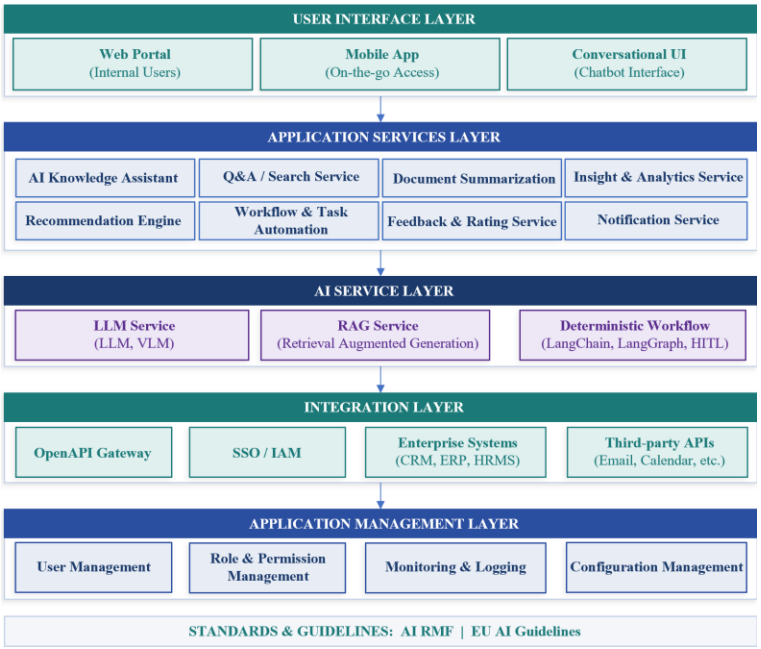


Figure 6. Application Architecture

The Technology Architecture complements these layers with an AI and ML models layer (LLM, embedding model, and VLM), a RAG architecture layer, the technology stack, an infrastructure layer (cloud platform, containerization, and CI/CD), and a security and operations layer (network security, encryption, backup and disaster recovery, and monitoring).

Prototype Development

A PoC prototype named TrustworthyAI was developed as the implementation artifact. Its knowledge base was created through an ingestion pipeline summarized in Table 5. The knowledge base was intentionally limited to company product knowledge documents as a quick-win use case, enabling sales personnel to obtain validated product information without relying on presales for every initial request, while presales remains essential for complex requirement analysis and solution design.

Table 5. Knowledge Base Creation Process

Stage	Process	Strategy	Output
1. Parsing	Documents are uploaded and converted into machine-readable content	Docling converts PDF documents into Markdown while preserving text and tables	Normalized Markdown content
2. Chunking	Content is divided into smaller knowledge units	Hybrid chunking following Markdown structure, preserving tables, and classifying by presales category; ~1,024 tokens with 100-token overlap	Structured chunks with metadata
3. Embedding	Each chunk is converted into a vector representation	Embedding model generates 1,024-dimensional vectors	Vector embeddings
4. Storage	Documents, chunks, embeddings, and metadata are stored	Supabase PostgreSQL and pgvector	Searchable organizational knowledge base
5. Retrieval	Relevant knowledge is retrieved for a user query	Hybrid vector and full-text search followed by reranking before passing context to the LLM	Relevant knowledge with source references

The prompt design follows the routing workflow described by Schluntz and Zhang (2024). Each request is first assigned a routing decision (needs_retrieval, needs_memory, direct_answer, or needs_clarification) and then, where retrieval is required, an intent: product_qa, pricing, tender_analysis, solution_arch, summarize_doc, compare_doc, product_mapping, risk_analysis, swot_bant, or general. Prompts are structured with XML tags that separate instructions, output schema, routing definitions, conversation history, and user input, and the routing prompt returns structured JSON containing the routing decision, intent, and a one-sentence reason. Table 6 summarizes the implemented technology stack, and Figure 7 shows the retrieval workflow.

Table 6. Technology Stack Summary

Layer	Technology	Purpose
Frontend	React	User interface and conversation management
Backend	FastAPI	API services and application logic
Workflow Orchestration	LangGraph	Workflow orchestration and RAG pipeline management
LLM Inference	Ollama	Local AI model deployment and inference
Large Language Model	Gemma3:4B	Question answering and response generation
Embedding Model	Qwen3-Embedding (1024-dim)	Query and chunk embedding
Vision Language Model	Qwen3-VL	OCR and multimodal document understanding
Document Processing	Docling	PDF and DOCX content extraction
Vector + full-text DB	Supabase (PostgreSQL + pgvector)	Structured knowledge storage
Conversation Storage	PostgreSQL	Chat history and session management

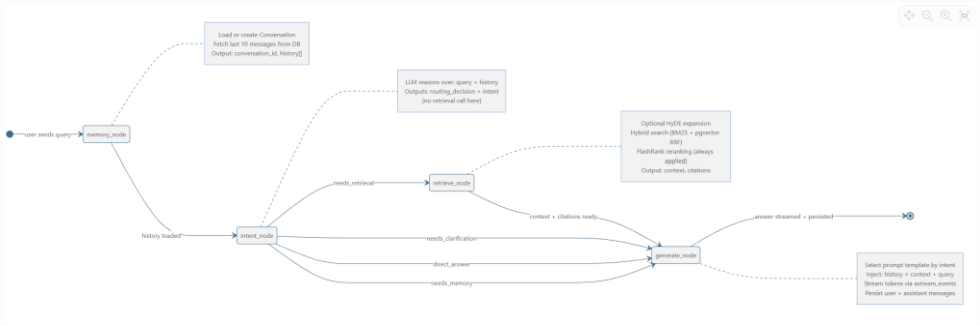


Figure 7. AI-Assisted Knowledge Retrieval Workflow

Pilot Implementation and Trustworthiness Controls

The pilot organized representative presales knowledge into eight categories (customer profile, opportunity information, solution documents, commercial documents, technical knowledge, organizational knowledge, project knowledge, and governance documents). Each asset carries standardized metadata such as category, industry, solution domain, owner, reviewer, validation status (Draft, Under Review, Approved, Archived), version, and source reference. The demonstration confirmed that the prototype retrieved relevant knowledge and generated grounded responses with source references, reducing manual searches across multiple repositories through a single conversational interface. Documents are uploaded, parsed, and checked using quality indicators such as OCR confidence and table extraction counts before indexing, and every generated answer is presented together with its supporting retrieved chunks and source references (Figure 8). Role-based access control restricts users to authorized knowledge, while knowledge owners review and validate content before it becomes available for AI retrieval.

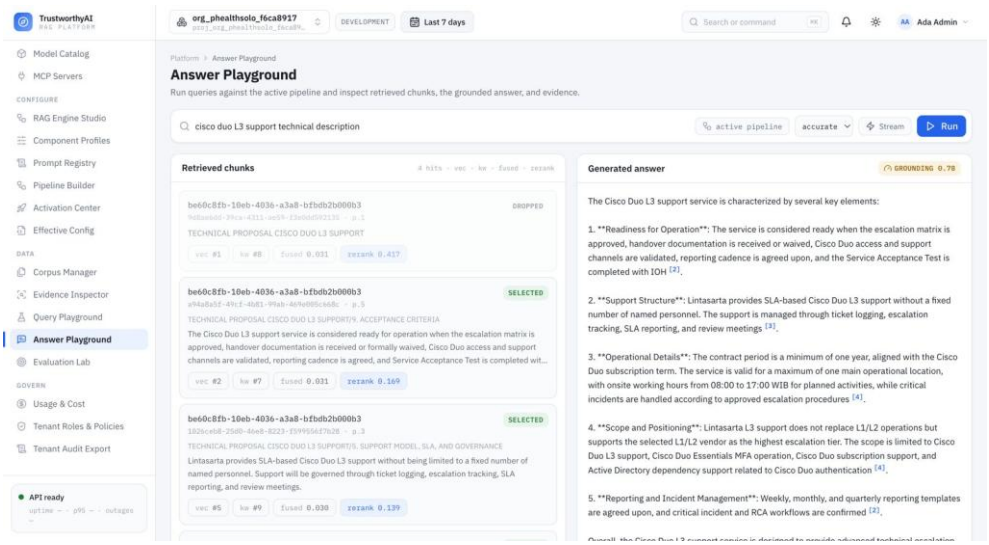


Figure 8. Source Traceability and Transparency

Trustworthiness Evaluation and Expert Validation

Six experts assessed the framework using fifteen statements across the three dimensions. The evaluation indicators were developed by translating the principles of the selected frameworks into assessment statements: A1 and A4 derive from TOGAF, A2 and A3 evaluate RAG-based retrieval and its trustworthiness, and A5 links the architecture to sales readiness, while G1–G5 derive from COBIT, the KM Lifecycle, NIST AI RMF, EU Guidelines, ISO/IEC 27001, and the PDP Law. The results are presented in Tables 9 to 10; all statements fall within the Highly Achieved category (4.21–5.00).

Table 7. Knowledge Structure Evaluation Results

Code	Assessment Statement	EV1	EV2	EV3	EV4	EV5	EV6	Mean	Weighted Mean
Weight		1.25	1.25	1.25	1.25	1.00	1.00		
S1	The proposed framework provides a clear structure for organizing organizational knowledge.	5	5	4	5	4	5	4.67	4.68
S2	The framework improves the consistency of early-stage sales preparation.	4	5	5	4	4	5	4.50	4.50
S3	The framework facilitates effective knowledge sharing and reuse across different business opportunities.	4	4	4	5	4	5	4.33	4.32
S4	The framework reduces dependency on individual knowledge and experience.	5	4	4	5	4	5	4.50	4.50
S5	The proposed framework supports continuous improvement through systematic knowledge capture and reuse.	5	5	4	5	4	5	4.67	4.68
Average								**4.53**	**4.54**

The Knowledge Structure dimension achieved a weighted mean of 4.54 (Highly Achieved). S1 and S5 obtained the highest scores, indicating strong agreement that the framework provides a clear knowledge structure and supports continuous improvement, while S3 received the lowest score, suggesting that knowledge classification and organizational adoption could further enhance cross-unit knowledge sharing. Experts also recommended strengthening metadata management and periodic knowledge reviews.

Table 8. Architecture Evaluation Results

Code	Assessment Statement	EV1	EV2	EV3	EV4	EV5	EV6	Mean	Weighted Mean
Weight		1.25	1.25	1.25	1.25	1.00	1.00		
A1	The proposed enterprise architecture appropriately integrates business, information, application, and technology components.	5	5	4	5	4	5	4.67	4.68
A2	The Retrieval-Augmented Generation (RAG) approach improves access to organizational knowledge.	5	5	5	5	4	5	4.83	4.86
A3	AI-generated responses are relevant, reliable, and traceable to organizational knowledge sources.	4	5	4	5	4	5	4.50	4.50
A4	The proposed architecture is suitable for enterprise implementation and future scalability.	4	4	4	5	4	5	4.33	4.32
A5	The proposed architecture effectively supports sales readiness during early-stage B2B engagement.	5	5	4	5	4	5	4.67	4.68
Average								**4.60**	**4.61**

The Architecture dimension achieved a weighted mean of 4.61 (Highly Achieved). A2 obtained the highest score, showing strong agreement that RAG improves access to organizational knowledge while maintaining contextual relevance. A4 received the lowest score, indicating that system integration, infrastructure readiness, and scalability planning should be addressed during full-scale deployment, including integration with CRM and document management systems.

Table 9. Governance Evaluation Results

Code	Assessment Statement	EV1	EV2	EV3	EV4	EV5	EV6	Mean	Weighted Mean
Weight		1.25	1.25	1.25	1.25	1.00	1.00		
G1	The framework clearly defines knowledge ownership and accountability.	5	5	4	5	5	4	4.67	4.68
G2	The framework provides appropriate knowledge validation and review mechanisms.	5	4	4	5	5	4	4.50	4.50
G3	The governance mechanisms improve the trustworthiness of AI-generated knowledge.	5	5	4	5	4	5	4.67	4.68
G4	The framework adequately addresses information security and personal data protection requirements.	4	4	4	5	5	4	4.33	4.32
G5	The proposed governance mechanisms support sustainable and trustworthy knowledge management.	5	5	5	5	4	5	4.83	4.86
Average								**4.60**	**4.61**

The Governance dimension also achieved a weighted mean of 4.61 (Highly Achieved). G5 obtained the highest score, indicating that the governance mechanisms provide a sustainable foundation for trustworthy knowledge management, while G4 received the lowest, suggesting that additional organizational policies and implementation guidelines on security and personal data protection would strengthen enterprise adoption. The aggregated results are shown in Table 10.

Table 10. Overall Trustworthiness Evaluation Results

Dimension	Mean	Weighted Mean	Success Criteria	Interpretation
Knowledge Structure	4.53	4.54	≥ 4.00	Highly Achieved
Architecture	4.60	4.61	≥ 4.00	Highly Achieved
Governance	4.60	4.61	≥ 4.00	Highly Achieved
Overall Average	**4.58**	**4.59**	**≥ 4.00**	**Highly Achieved**

With an overall mean of 4.58 and a weighted mean of 4.59, all dimensions exceeded the 4.00 threshold, and the PoC demonstrated the intended end-to-end workflow; therefore, the success criteria were met and hypotheses H1–H3 are supported within the scope of expert validation. The qualitative comments did not identify fundamental weaknesses requiring redesign. Instead, the experts emphasized enterprise integration and scalability (EV1), measurable business outcomes and adoption (EV2), the continuing role of presales for complex solutions (EV3), taxonomy and periodic review (EV4), clear ownership and governance reviews (EV5), and continuous monitoring of retrieval quality (EV6), as summarized in Table 11.

Table 11. Summary of Expert Recommendations

Evaluation Dimension	Recommendation
Knowledge Structure	Establish standardized taxonomy and metadata; conduct periodic knowledge reviews to maintain accuracy and relevance.

Evaluation Dimension	Recommendation
Architecture	Integrate with CRM and document management systems; monitor retrieval quality and optimize indexing as the repository grows.
Governance	Define clear knowledge ownership and review responsibilities; conduct periodic governance and compliance reviews.
Business Adoption	Integrate the framework into daily sales and presales activities and link implementation to measurable sales-readiness outcomes.
AI Trustworthiness	Maintain source traceability, retrieval monitoring, controlled access, and human oversight as the system evolves.

The findings indicate that the primary challenge affecting sales readiness is not a lack of organizational knowledge, but the lack of a structured approach to managing and reusing it. The framework addresses this by integrating structured knowledge management, a centralized repository, governance mechanisms, and AI-assisted retrieval, demonstrating that trustworthy AI depends on high-quality organizational knowledge supported by effective governance and human oversight. This is consistent with the view that AI supports information-intensive sales activities while complementing human capabilities, and it extends technically oriented AI-KM implementations (Wen et al., 2025) by embedding governance and lifecycle controls into the architecture.

Five critical success factors were identified: management commitment, knowledge quality through validation and regular updates, user adoption through integration into daily sales and presales activities, measurable sales-readiness outcomes (e.g., reducing preparation time from approximately one week to one working day and increasing opportunity creation by approximately 30% year-on-year), and governance through clear ownership and review. Key implementation challenges are incomplete or outdated knowledge, users' willingness to contribute, continuous validation effort, and integration with existing enterprise systems. The findings are limited to a single-organization PoC evaluated mainly through qualitative expert assessment, without comparison of AI models or RAG techniques and without long-term adoption or operational measurement.

CONCLUSION

This study developed and evaluated a trustworthy AI-enabled knowledge management framework to support early-stage sales readiness in business-to-business (B2B) technology service organizations using the Design Science Research (DSR) methodology. The findings identified knowledge fragmentation, the absence of a centralized knowledge repository, and inadequate governance mechanisms as key challenges associated with inconsistent sales readiness. To address these challenges, the proposed framework integrated the Socialization, Externalization, Combination, and Internalization (SECI) model; the Knowledge Management Lifecycle; The Open Group Architecture Framework (TOGAF); retrieval-augmented generation (RAG); and governance principles derived from Control Objectives for Information and Related Technologies (COBIT), the NIST AI Risk Management Framework (AI RMF), the European Union's Ethics Guidelines for Trustworthy AI, ISO/IEC 27001, and Indonesia's Personal Data Protection Law. The framework was demonstrated through the TrustworthyAI proof-of-concept (PoC) prototype and evaluated by six multidisciplinary experts. The evaluation produced experience-weighted mean scores of 4.54 for Knowledge Structure, 4.61 for Architecture, 4.61 for Governance, and 4.59 overall, all classified as Highly Achieved according to the study's evaluation criteria. These results indicated favorable expert assessments across the three evaluated dimensions and supported

the achievement of the study's design and evaluation objectives. The findings suggested that improving sales readiness required not only AI capabilities but also structured organizational knowledge, enterprise architecture alignment, effective governance, and human oversight. Accordingly, organizations should prioritize establishing centralized and regularly maintained knowledge repositories, assigning clear knowledge ownership and governance responsibilities, implementing knowledge validation procedures before AI-assisted retrieval, introducing AI capabilities incrementally through pilot projects, and continuously monitoring knowledge quality. Future research should evaluate the framework across different organizations and industries, examine its integration with customer relationship management (CRM), enterprise resource planning (ERP), and document management systems, and assess production readiness in terms of scalability, security, and operational monitoring. Longitudinal studies should also investigate measurable business outcomes, including sales preparation time, opportunity generation, knowledge reuse, and sales win rates. Further research could explore multi-agent AI systems, knowledge graphs, and advanced RAG techniques while maintaining transparency, accountability, and human oversight.

FUNDING

This study did not receive any specific funding from governmental institutions, private organizations, or other funding agencies.

CONFLICT OF INTEREST

The authors declare that there are no conflicts of interest related to the conduct of the research, preparation, or publication of this article.

AUTHOR CONTRIBUTIONS

All authors contributed to the preparation of this article. The first author, Gerry Gusman Wibowo, was responsible for the conceptualization of the study, data collection and analysis, interpretation of the results, and manuscript preparation. The other authors contributed to the review, editing, and final approval of the manuscript.

GENERATIVE AI STATEMENT

The authors declare that no generative artificial intelligence was used in the preparation of this article. (Belhadi & Benkhellil, 2026)

REFERENCES

Abdelaziz, M., & Khalil, T. (2023a). Exploring the presales process in high-tech companies. *IEEE Engineering Management Review*, *51*(3), 51–57.

Abdelaziz, M., & Khalil, T. (2023b). A knowledge management framework for presales operations in high-technology companies. In *2023 International Conference on IT Innovation and Knowledge Discovery (ITIKD)* (pp. 1–6). IEEE.

Agbabiaka, O., Ojo, A., & Connolly, N. (2025). Requirements for trustworthy AI-enabled automated decision-making in the public sector: A systematic review. *Technological Forecasting and Social Change*, *215*, 124076. https://doi.org/10.1016/j.techfore.2025.124076

Alawamleh, M., Shammas, N., Alawamleh, K., & Bani Ismail, L. (2024). Examining the limitations of AI in business and the need for human insights using interpretive structural

modelling. *Journal of Open Innovation: Technology, Market, and Complexity*, *10*(3), 100338. https://doi.org/10.1016/j.joitmc.2024.100338

Ali, M., Khan, T. I., Khattak, M. N., & Şener, İ. (2024). Synergizing AI and business: Maximizing innovation, creativity, decision precision, and operational efficiency in high-tech enterprises. *Journal of Open Innovation: Technology, Market, and Complexity*, *10*(3), 100352. https://doi.org/10.1016/j.joitmc.2024.100352

Andersson, S., Aagerup, U., Svensson, L., & Eriksson, S. (2024). Challenges and opportunities in the digitalization of the B2B customer journey. *Journal of Business & Industrial Marketing*, *39*(13), 160–174. https://doi.org/10.1108/JBIM-12-2023-0714

Clemen, R. T., & Winkler, R. L. (1999). Combining probability distributions from experts in risk analysis. *Risk Analysis*, *19*(2), 187–203. https://doi.org/10.1111/j.1539-6924.1999.tb00399.x

Colson, A. R., & Cooke, R. M. (2018). Expert elicitation: Using the classical model to validate experts' judgments. *Review of Environmental Economics and Policy*, *12*(1), 113–132. https://doi.org/10.1093/reep/rex022

Cooke, R. M. (1991). *Experts in uncertainty: Opinion and subjective probability in science*. Oxford University Press.

De Keyser, A., Antonetti, P., Rouziou, M., Béal, M., Wang, Z.-H., Grégoire, Y., & Lussier, B. (2025). Understanding the B2B customer experience and journey: A convergence-based lens. *Journal of Business Research*, *198*, 115481. https://doi.org/10.1016/j.jbusres.2025.115481

Gartner. (2024). *2024 global software buying trends*. Gartner Digital Markets.

Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research. *MIS Quarterly*, *28*(1), 75–105. https://doi.org/10.2307/25148625

Huang, M.-H., & Rust, R. T. (2021). Artificial intelligence in service. *Journal of Service Research*, *24*(1), 3–19. https://doi.org/10.1177/1094670520902266

James, A., Trovati, M., & Bolton, S. (2025). Retrieval-augmented generation to generate knowledge assets and creation of action drivers. *Applied Sciences*, *15*(11), 6247. https://doi.org/10.3390/app15116247

Johnston, W. J., & Cortez, R. M. (2024). Business-to-business digitalization, artificial intelligence, and social action. *Journal of Business Research*, *172*, 113952. https://doi.org/10.1016/j.jbusres.2023.113952

Krueger, R. A., & Casey, M. A. (2015). *Focus groups: A practical guide for applied research (5th ed.)*. SAGE.

Lundin, L., & Kindström, D. (2023). Digitalizing customer journeys in B2B markets. *Journal of Business Research*, *157*, 113639. https://doi.org/10.1016/j.jbusres.2022.113639

Mittal, V., Astvansh, V., & Singh, S. (2025). Effective B2B sales growth strategy: Evidence from a longitudinal study of 44,000 B2B customers. *Management and Business Review*.

Schluntz, E., & Zhang, B. (2024, December 19). *Building effective agents*. Anthropic. https://www.anthropic.com/engineering/building-effective-agents

Uren, V., & Edwards, J. S. (2023). Technology readiness and the organizational journey towards AI adoption: An empirical study. *International Journal of Information Management*, *68*, 102588. https://doi.org/10.1016/j.ijinfomgt.2022.102588

vom Brocke, J., Hevner, A., & Maedche, A. (2020). Introduction to design science research. In J. vom Brocke, A. Hevner, & A. Maedche (Eds.), *Design science research. Cases* (pp. 1–13). Springer. https://doi.org/10.1007/978-3-030-46781-4_1

Webster, F. E., Jr., & Wind, Y. (1972). A general model for understanding organizational buying behavior. *Journal of Marketing*, *36*(2), 12–19. https://doi.org/10.1177/002224297203600204

Wen, H., Wang, S., Jiang, C., Zhang, F., Li, J., & Luo, X. (2025). AI–KM: Knowledge enhancement with RAG and workflow. *SoftwareX*, *31*, 102349.

Wiersema, F. (2013). The B2B agenda: The current state of B2B marketing and a look ahead. *Industrial Marketing Management*, *42*(4), 470–488. https://doi.org/10.1016/j.indmarman.2013.02.015