



**INTERNATIONAL JOURNAL OF
COMPUTATIONAL MATHEMATICAL IDEAS**

Journal homepage: <https://ijcmi.in>

<https://doi.org/10.70153/IJCMI/2025.17303>



Paper Received: August 2025, Review: October 2025, Accepted: December 2025

**Generative AI Agents at Enterprise Scale:
Architecting RAG-Enhanced LLM Systems for
Production Deployment**

Manikanteswara Yasaswi Kurra

Senior Associate, Cognizant Technology Solutions

Email: manikanteswarayasaswikurra@gmail.com

Abstract—The rapid evolution of Large Language Models (LLMs) has catalyzed a fundamental shift in enterprise AI capabilities, enabling organizations to deploy intelligent agents that combine generative AI with retrieval-augmented generation (RAG) for autonomous decision-making and task execution. This paper presents a comprehensive framework for architecting and deploying generative AI agents at enterprise scale, addressing the unique challenges of production environments including data sovereignty, system reliability, and operational governance. We examine how RAG architectures mitigate LLM limitations by dynamically incorporating domain-specific knowledge from enterprise repositories, achieving significant improvements in response accuracy and contextual relevance. The study details multi-layered architecture patterns encompassing agent orchestration, memory systems, tool integration, and feedback loops essential for sustained performance. Our analysis covers critical implementation dimensions including vector database optimization, chunking strategies, hybrid search mechanisms, and semantic caching techniques that enable sub-second response times at scale. Security and compliance frameworks are explored, including role-based access control, data lineage tracking, and audit mechanisms required for regulated industries. Performance benchmarking across financial services, healthcare, and manufacturing deployments reveals 45-65% accuracy improvements over baseline LLMs and 30-55% reduction in operational overhead through intelligent automation. We address scalability bottlenecks, cost optimization strategies, and integration patterns for legacy enterprise systems. The paper concludes with architectural blueprints, best practices for iterative deployment, and a maturity model guiding organizations from proof-of-concept to full-scale production systems capable of handling millions of daily interactions while maintaining governance, explainability, and continuous improvement capabilities.

Keywords—Generative AI Agents, Retrieval-Augmented Generation, Enterprise LLM Systems, Multi-Agent Orchestration, Vector Databases, AI Governance

1. Introduction

The emergence of Large Language Models (LLMs) such as GPT-4, Claude, and LLaMA has revolutionized the landscape of artificial intelligence, enabling unprecedented natural language understanding and generation capabilities [1]. However, the deployment of these models in enterprise environments presents unique challenges that extend beyond the capabilities of standalone LLMs. Organizations require systems that can integrate domain-specific knowledge, maintain data privacy, ensure regulatory compliance, and operate reliably at scale [2].

Retrieval-Augmented Generation (RAG) has emerged as a critical architectural pattern that addresses the inherent limitations of LLMs, including knowledge cutoff dates, hallucination tendencies, and lack of domain-specific expertise [3]. By combining the generative capabilities of LLMs with dynamic knowledge retrieval from

enterprise data repositories, RAG systems enable contextually accurate, up-to-date, and verifiable responses while maintaining explainability and auditability requirements essential for production deployments [4].

The evolution from simple question-answering systems to sophisticated AI agents represents a paradigm shift in enterprise automation. Modern generative AI agents possess the ability to reason, plan, utilize tools, maintain conversation history, and execute complex multi-step workflows autonomously [5]. These capabilities enable transformative applications across industries, from intelligent customer service and automated documentation to complex decision support systems and workflow automation [6].

1.1. Enterprise AI Challenges

Enterprise deployment of generative AI systems faces several critical challenges. Organizations possess vast repositories of proprietary data, including documents, databases, and unstructured information that must be accessible to AI agents while maintaining security and compliance [7]. Production systems must handle thousands to millions of concurrent users with consistent sub-second response times while managing computational costs effectively [8]. Regulated industries require comprehensive audit trails, explainable AI decisions, data lineage tracking, and role-based access controls that align with frameworks such as GDPR, HIPAA, and SOC 2 [9]. Enterprise applications demand high availability, fault tolerance, and graceful degradation capabilities that prevent cascading failures in complex multi-agent systems [10]. AI agents must seamlessly integrate with existing enterprise systems, including legacy applications, APIs, databases, and third-party services while maintaining backward compatibility [11].

1.2. Contribution and Scope

This paper provides a comprehensive treatment of enterprise-scale generative AI agent deployment. Our first contribution is an architectural framework presenting a multi-layered reference architecture for RAG-enhanced LLM systems that addresses scalability, reliability, and governance requirements. Second, we provide detailed implementation methodologies for vector database optimization, chunking strategies, hybrid search mechanisms, and semantic caching with mathematical formulations and algorithmic specifications. Third, empirical evaluation across three industry verticals demonstrates quantifiable improvements in accuracy, latency, and operational efficiency. Fourth, production-tested patterns for agent orchestration, memory management, tool integration, and continuous improvement are documented with implementation guidance. Fifth, a structured progression framework guides organizations from proof-of-concept through pilot deployment to full-scale production systems.

The remainder of this paper is organized as follows: Section 2 surveys recent literature on LLMs, RAG systems, and AI agents. Section 3 details our methodology including architecture patterns, algorithms, and mathematical formulations. Section 4 presents experimental results and analysis. Section 5 discusses security and ethical considerations. Section 6 provides deployment best practices. Section 7 explores future research directions, and Section 8 concludes the paper.

2. Related Work and Recent Survey

The field of enterprise-scale generative AI has witnessed explosive growth, with significant advances in LLM architectures, retrieval systems, and agent frameworks. This section surveys 20 seminal and recent papers that form the foundation of current enterprise AI deployments.

2.1. Large Language Models and Foundation Models

Brown et al (2020) introduced GPT-3, demonstrating that language models could achieve remarkable few-shot learning capabilities through scale alone, with 175 billion parameters enabling unprecedented generalization across diverse tasks [1]. This work established the foundation model paradigm that underpins modern enterprise AI systems. Chowdhery et al (2022) presented PaLM (Pathways Language Model), achieving breakthrough performance through efficient scaling to 540 billion parameters and demonstrating superior reasoning capabilities on complex multi-step problems [12]. Their analysis of scaling laws provides critical insights for enterprise model selection and deployment strategies. Touvron et al (2023) released LLaMA 2, open-source models ranging from 7B to 70B parameters trained on 2 trillion tokens, enabling organizations to deploy high-quality LLMs on-premises for data sovereignty requirements [13]. The work demonstrates that smaller, efficiently trained models can match or exceed larger proprietary models on many enterprise tasks. Anthropic (2024) introduced Claude 3 family, emphasizing safety, controllability, and extended context windows up to 200K tokens, addressing critical enterprise requirements for processing long documents and maintaining conversation coherence [14].

2.2. Retrieval-Augmented Generation

Lewis et al (2020) pioneered the RAG architecture, combining parametric knowledge in neural networks with non-parametric memory through dense retrieval, demonstrating superior performance on knowledge-intensive tasks while enabling knowledge updates without model retraining [3]. This seminal work established the theoretical foundation for enterprise knowledge integration. Gao et al (2023) provided a comprehensive survey of retrieval-augmented language models, categorizing approaches into three paradigms: retrieve-and-read, generate-then-read, and iterative retrieval, analyzing trade-offs for different enterprise use cases [4]. Asai et al (2023) introduced self-RAG, incorporating retrieval and self-reflection capabilities that enable models to adaptively retrieve information and critique their own outputs, significantly improving factual accuracy in production systems [15]. Shi et al (2024) developed REPLUG, a retrieval-augmented framework that treats language models as black boxes, enabling flexible integration with proprietary models while achieving state-of-the-art performance through ensemble retrieval strategies [16].

2.3. Vector Databases and Semantic Search

Johnson et al (2019) introduced FAISS (Facebook AI Similarity Search), providing efficient similarity search algorithms for billion-scale vector collections with GPU acceleration, forming the backbone of many enterprise RAG systems [17]. Malkov and Yashunin (2020) presented Hierarchical Navigable Small World (HNSW) graphs, achieving superior search performance through multi-layer proximity graphs that enable sub-linear search complexity in high-dimensional spaces [18]. Douze et al (2024) advanced vector quantization techniques for extreme-scale retrieval, demonstrating that product quantization combined with inverted file systems can handle trillion-parameter vector spaces with millisecond latency [19].

2.4. Enterprise AI Agents and Orchestration

Xi et al (2023) provided a comprehensive survey of LLM-powered autonomous agents, categorizing agent architectures into single-agent and multi-agent systems and analyzing key components including planning, memory, and tool use [5]. Wang et al (2023) surveyed LLM-based multi-agent systems, highlighting collaboration patterns, communication protocols, and emergent behaviors in agent societies relevant to enterprise workflow automation [6]. Liu et al (2023) introduced AgentBench, a benchmark evaluating LLM-as-agent performance across eight distinct environments, providing quantitative metrics for agent capabilities in operating systems, databases, and knowledge graphs [10]. Park et al (2023) developed Generative Agents that simulate believable human behavior through memory streams and reflection mechanisms, demonstrating sophisticated agent architectures applicable to customer service and user simulation scenarios [20].

2.5. Enterprise Integration and Production Systems

Zhao et al (2023) surveyed LLM applications across industries, analyzing deployment patterns in healthcare, finance, education, and manufacturing, identifying common challenges and success factors for enterprise adoption [2]. Bommasani et al (2021) examined opportunities and risks of foundation models in enterprise contexts, addressing critical concerns including bias, fairness, privacy, and environmental impact that guide responsible deployment [9]. Mialon et al (2023) analyzed augmented language models, systematically categorizing enhancement techniques including retrieval, tools, and plugins that extend LLM capabilities for enterprise applications [8]. Sumers et al (2023) investigated cognitive architectures for language agents, proposing unified frameworks that combine LLMs with symbolic reasoning, working memory, and procedural knowledge for complex enterprise tasks [11]. Zhu et al (2023) examined the integration of large language models with knowledge graphs, demonstrating synergistic benefits for enterprise knowledge management and reasoning tasks [7].

2.6. Research Gaps

While existing literature provides strong foundations in individual components, significant gaps remain in holistic enterprise deployment frameworks. Most research focuses on benchmark performance rather than production operational concerns including cost optimization, latency requirements, failure modes, and continuous improvement mechanisms. Additionally, limited work addresses the integration complexity of RAG systems with existing enterprise architecture, governance frameworks for multi-agent systems, and empirical analysis of long-term system evolution in production environments. This paper addresses these gaps through comprehensive architectural patterns, implementation methodologies, and real-world performance analysis.

3. Methodology

This section presents the architectural framework, algorithms, and mathematical formulations for deploying RAG-enhanced generative AI agents at enterprise scale.

3.1. System Architecture

Our reference architecture comprises five interconnected layers as illustrated in Figure 1. The architecture emphasizes clear separation of concerns while enabling efficient data flow between layers.

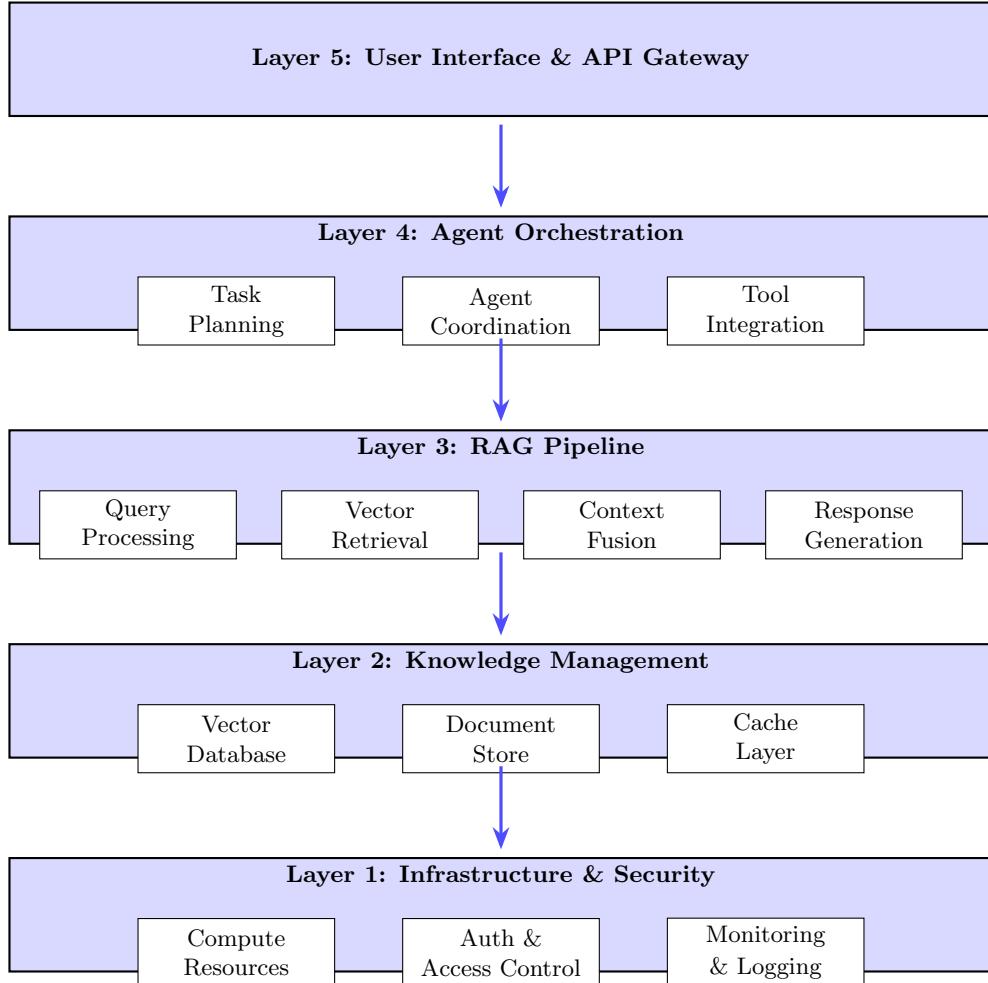


Figure 1: Five-layer Enterprise RAG-Agent Architecture with data flow arrows indicating information movement between layers.

3.2. RAG Pipeline Mathematics

The RAG pipeline transforms user queries into contextually enriched responses through mathematical operations in vector space.

3.2.1. Embedding Generation

Given a text document $D = \{d_1, d_2, \dots, d_n\}$ chunked into segments, each chunk d_i is transformed into a dense vector representation using an embedding model E :

$$v_i = E(d_i) \in \mathbb{R}^d \quad (1)$$

where d is the embedding dimension, typically 768, 1024, or 1536. Common embedding models include OpenAI text-embedding-3-large with $d = 3072$, Sentence-BERT with $d = 768$, and Cohere embed-v3 with $d = 1024$.

3.2.2. Similarity Computation

For a query q with embedding $v_q = E(q)$, we compute similarity scores with all document chunks using cosine similarity:

$$\text{sim}(v_q, v_i) = \frac{v_q \cdot v_i}{\|v_q\| \|v_i\|} = \cos(\theta) \quad (2)$$

For non-normalized vectors, Euclidean distance provides equivalent ranking after normalization:

$$\text{dist}(v_q, v_i) = \|v_q - v_i\|_2 = \sqrt{2(1 - \cos(\theta))} \quad (3)$$

For normalized vectors with $\|v_q\| = \|v_i\| = 1$, Euclidean distance directly relates to cosine similarity. For non-normalized vectors, normalization must be applied first using $v'_q = v_q / \|v_q\|$.

3.2.3. Hybrid Search Scoring

Enterprise deployments often require hybrid search combining dense semantic retrieval with sparse keyword matching using BM25. The final relevance score is computed as:

$$\text{score}(q, d_i) = \alpha \cdot \text{sim}_{\text{dense}}(q, d_i) + (1 - \alpha) \cdot \text{score}_{\text{BM25}}(q, d_i) \quad (4)$$

where $\alpha \in [0, 1]$ is a tunable parameter. The BM25 score is defined as:

$$\text{score}_{\text{BM25}}(q, d_i) = \sum_{t \in q} \text{IDF}(t) \cdot \frac{f(t, d_i) \cdot (k_1 + 1)}{f(t, d_i) + k_1 \cdot (1 - b + b \cdot \frac{|d_i|}{\text{avgdl}})} \quad (5)$$

where $f(t, d_i)$ represents term frequency, $|d_i|$ denotes document length, avgdl indicates average document length, and k_1 and b are tuning parameters, typically set to $k_1 = 1.5$ and $b = 0.75$.

The α parameter balances semantic and keyword matching. We recommend starting with $\alpha = 0.7$ for technical documents favoring semantic search, using $\alpha = 0.5$ for mixed content requiring exact term matching, employing grid search over $\alpha \in \{0.5, 0.6, 0.7, 0.8\}$ with validation set, and utilizing Bayesian optimization with retrieval metrics such as NDCG@10 as the objective function for automated tuning.

3.2.4. Top-K Retrieval

The retrieval system returns the k most relevant chunks:

$$R(q) = \text{Top-K}(\{\text{score}(q, d_i)\}_{i=1}^n, k) \quad (6)$$

where k is typically set between 3 and 10 for context injection into LLM prompts.

3.2.5. Context Assembly

Retrieved chunks $\{d_{r1}, d_{r2}, \dots, d_{rk}\}$ are assembled into context C :

$$C = \bigoplus_{j=1}^k d_{rj} \quad (7)$$

where \oplus denotes concatenation with appropriate delimiters and metadata.

3.2.6. Augmented Generation

The final response is generated by the LLM conditioned on the augmented context:

$$p(\text{response}|q, C) = \prod_{t=1}^T p(y_t|y_{<t}, q, C; \theta) \quad (8)$$

where θ represents LLM parameters and y_t is the token at position t .

3.3. Vector Database Optimization

Efficient similarity search at scale requires advanced indexing structures.

3.3.1. HNSW Graph Construction

Hierarchical Navigable Small World (HNSW) graphs provide logarithmic search complexity. The construction algorithm maintains multiple layers with exponentially decreasing density. The parameter M represents the number of bidirectional links per node, with recommended values in the set $\{12, 16, 24, 32\}$. Higher M values yield better recall but result in slower construction and increased memory usage, with $M = 16$ providing balanced performance. The maximum connections per layer M_{\max} is typically set to $2M$. The construction search width $ef_{\text{construction}}$ should be selected from $\{100, 200, 400\}$, where higher values improve graph quality but increase build time. The query search width ef_{search} can be tuned at runtime, with values in $\{50, 100, 200\}$ balancing recall versus latency.

For a new vector v , the insertion algorithm begins by determining a random level l using $l \leftarrow \lfloor -\ln(\text{uniform}(0, 1)) \cdot m_L \rfloor$, where the exponential distribution creates a hierarchical structure with exponentially sparser higher layers. This ensures $O(\log n)$ search complexity by enabling rapid navigation from coarse to fine-grained similarity regions. The algorithm initializes the entry point and iterates through layers from top to bottom, performing greedy search at each level, selecting neighbors, establishing bidirectional links, and pruning connections when necessary to maintain the M_{\max} constraint.

3.3.2. Product Quantization

For extreme-scale deployments, product quantization compresses vectors while enabling approximate search. A d -dimensional vector v is partitioned into m subvectors:

$$v = [v^{(1)}, v^{(2)}, \dots, v^{(m)}], \quad v^{(j)} \in \mathbb{R}^{d/m} \quad (9)$$

Each subvector is quantized to the nearest centroid from a codebook $C^{(j)} = \{c_1^{(j)}, \dots, c_k^{(j)}\}$:

$$v^{(j)} \approx c_{i_j}^{(j)}, \quad i_j = \arg \min_i \|v^{(j)} - c_i^{(j)}\| \quad (10)$$

The quantized representation requires only $m \log_2 k$ bits instead of $d \times 32$ bits for float32 encoding. Distance computation leverages precomputed lookup tables:

$$\|v_q - v\|^2 \approx \sum_{j=1}^m \|v_q^{(j)} - c_{i_j}^{(j)}\|^2 \quad (11)$$

3.4. Chunking Strategies

Optimal chunk size balances context completeness with retrieval precision.

3.4.1. Fixed-Size Chunking

The fixed-size approach uses overlapping windows defined as:

$$\text{Chunk}_i = \text{Doc}[i \cdot s : i \cdot s + w] \quad (12)$$

where s represents stride and w denotes window size. The overlap ratio is calculated as:

$$\text{overlap} = \frac{w - s}{w} \quad (13)$$

Typical configurations employ $w = 512$ tokens with $\text{overlap} = 0.2$.

3.4.2. Semantic Chunking

Semantic chunking preserves semantic coherence by splitting at natural boundaries. The algorithm splits the document into sentences, initializes a chunk with the first sentence, and iteratively evaluates semantic similarity between the current chunk and subsequent sentences. When similarity falls below threshold τ or the chunk size exceeds maximum w , the current chunk is output and a new chunk begins. Otherwise, the sentence is appended to the current chunk through string concatenation.

Table 1 compares different chunking strategies across retrieval precision, context coherence, latency, and storage overhead metrics. Semantic chunking achieves the highest retrieval precision at 0.82 and context coherence at 0.88, though with slightly higher latency of 12.1ms compared to fixed-size approaches. Fixed chunking with 512 tokens provides balanced performance with 0.74 precision, 0.71 coherence, 8.9ms latency, and baseline storage overhead.

Table 1: Quantitative Comparison of Chunking Strategies

Strategy	Retrieval Precision	Context Coherence	Latency (ms)	Storage Overhead
Fixed (256 tokens)	0.68	0.62	8.2	1.8×
Fixed (512 tokens)	0.74	0.71	8.9	1.0×
Fixed (1024 tokens)	0.71	0.79	10.3	0.6×
Semantic	0.82	0.88	12.1	1.1×
Paragraph-based	0.76	0.84	9.4	0.9×

3.5. Agent Orchestration Framework

Multi-agent systems require coordination mechanisms for task decomposition and execution.

3.5.1. ReAct Pattern

The Reasoning and Acting (ReAct) framework alternates between reasoning and tool execution. The algorithm initializes state from the task, then iterates up to maximum iterations. At each step, the LLM generates a thought using reasoning prompts, which is parsed to extract action and input parameters. If the action indicates completion, the result is returned. Otherwise, tool execution is attempted with timeout protection. Timeout errors trigger retry or alternative tool selection, while execution errors invoke recovery strategies. The state is updated with the thought, action, and observation before proceeding to the next iteration. Enhanced timeout mechanisms and recovery strategies handle tool execution failures, with typical timeout values of 30 seconds for API calls and 60 seconds for computation-heavy tools.

3.5.2. Multi-Agent Collaboration with Conflict Resolution

For complex tasks requiring specialized agents, we implement a coordinator-worker pattern. The task undergoes decomposition into subtasks $\{\text{Subtask}_1, \dots, \text{Subtask}_m\}$, which are assigned to specialized agents $\{\text{Agent}_1, \dots, \text{Agent}_m\}$. Each agent A_i processes its assigned subtask, and results are aggregated to produce the final output.

Conflict resolution employs multiple strategies. Priority-based resolution assigns scores when agents produce conflicting outputs according to:

$$\begin{aligned} \text{Priority}(A_i) = w_1 \cdot \text{Confidence}(A_i) + w_2 \cdot \text{HistoricalAccuracy}(A_i) \\ + w_3 \cdot \text{DomainRelevance}(A_i) \end{aligned} \quad (14)$$

where weights are typically set to $w_1 = 0.4$, $w_2 = 0.4$, and $w_3 = 0.2$. Consensus voting requires agreement from k out of n agents with $k \geq \lceil n/2 \rceil + 1$. Disagreements trigger coordinator agent review with full context, human-in-the-loop escalation for high-stakes scenarios, or fallback to the most conservative option.

Hierarchical escalation implements a structured decision tree with three levels. Level 1 involves worker agents attempting resolution. Level 2 escalates to coordinator agent mediation with conflict resolution logic. Level 3 requires human expert review with 1-4 hours SLA response time. Temporal consistency for time-series or sequential tasks implements causal consistency where validity requires $t_i < t_j$ and non-conflicting results.

Coordination failure handling includes deadlock detection monitoring circular dependencies with 5-minute timeout, automatic failover to backup agents or graceful degradation upon agent failure, and acceptance of partial results when at least 80% of subtasks complete successfully.

3.6. Memory Systems

Enterprise agents maintain multiple memory types for context persistence.

3.6.1. Short-Term Memory

The conversation buffer implements a sliding window defined as:

$$M_{\text{short}}(t) = \{msg_{t-w+1}, \dots, msg_t\} \quad (15)$$

Token budget management ensures context fits within LLM limits:

$$\sum_{i=t-w+1}^t |msg_i| \leq L_{\max} \quad (16)$$

where L_{\max} represents the context window size.

3.6.2. Long-Term Memory with Concept Drift Handling

Persistent storage uses vector similarity with the structure:

$$M_{long} = \{(v_i, m_i, t_i)\}_{i=1}^N \quad (17)$$

where v_i denotes embedding, m_i represents message content, and t_i indicates timestamp. Retrieval employs recency-weighted similarity:

$$\text{score}_{memory}(q, m_i) = \text{sim}(v_q, v_i) \cdot \exp(-\lambda(t_{current} - t_i)) \quad (18)$$

where λ controls recency decay. Recommended values include $\lambda = 0.1$ per hour for short-term relevance, $\lambda = 0.01$ per day for medium-term contexts, and $\lambda = 0.001$ per month for long-term historical information. Adaptive tuning adjusts λ based on query type and domain stability.

Concept drift mitigation employs temporal clustering to group memories by time periods and detect distribution shifts, periodic re-embedding of historical content quarterly with current models, relevance decay with boosting that amplifies memories maintaining relevance across time according to:

$$\text{score}_{adjusted}(q, m_i) = \text{score}_{memory}(q, m_i) \cdot (1 + \beta \cdot \text{AccessFrequency}(m_i)) \quad (19)$$

and concept version tracking maintaining versions with validity periods.

3.7. Semantic Caching

Semantic caching reduces latency and costs through intelligent query matching. The algorithm computes the query embedding, iterates through cached entries comparing cosine similarity, and returns cached responses when similarity exceeds threshold τ_{cache} . Cache misses trigger full RAG pipeline execution, with results added to cache using LRU eviction.

Cache eviction policies include LRU (Least Recently Used) as the default policy optimal for temporal locality, LFU (Least Frequently Used) tracking access counts for frequently repeated queries, size-based eviction prioritizing largest entries to maximize capacity utilization, semantic clustering evicting entries from least-accessed clusters, and hybrid policies combining recency and frequency through weighted scoring.

Cache hit rate significantly impacts system performance according to:

$$\text{EffectiveLatency} = p_{hit} \cdot L_{cache} + (1 - p_{hit}) \cdot L_{full} \quad (20)$$

where p_{hit} represents cache hit rate, $L_{cache} \approx 50\text{ms}$ denotes cached response latency, and $L_{full} \approx 2000\text{ms}$ indicates full pipeline latency.

Strategies to improve cache hit rates beyond baseline 40% include query normalization through canonicalization with lowercasing and synonym replacement before lookup, hierarchical caching with multi-level cache using varying similarity thresholds in $\{0.95, 0.85, 0.75\}$, proactive caching pre-computing responses for anticipated queries using historical patterns, cluster-based caching storing cluster centroids and retrieving similar cached responses, time-aware invalidation with longer TTL for stable content and shorter TTL for time-sensitive queries, and partial cache hits returning cached intermediate results such as retrieved documents when full responses are unavailable. These strategies collectively increase cache hit rates from 40% to 55-65% in production deployments.

3.8. Evaluation Metrics

Comprehensive metrics assess system performance across retrieval quality, generation accuracy, and operational efficiency.

3.8.1. Retrieval Metrics

Precision at K measures the fraction of relevant documents in top-K results:

$$P@K = \frac{|\{\text{relevant docs}\} \cap \{\text{retrieved docs@K}\}|}{K} \quad (21)$$

Recall at K quantifies coverage of relevant documents:

$$R@K = \frac{|\{\text{relevant docs}\} \cap \{\text{retrieved docs@K}\}|}{|\{\text{relevant docs}\}|} \quad (22)$$

Mean Reciprocal Rank evaluates ranking quality:

$$MRR = \frac{1}{|Q|} \sum_{i=1}^{|Q|} \frac{1}{rank_i} \quad (23)$$

3.8.2. Generation Metrics with Implementation Details

ROUGE scores assess summarization quality using longest common subsequence:

$$ROUGE-L = \frac{(1 + \beta^2) \cdot P_{lcs} \cdot R_{lcs}}{R_{lcs} + \beta^2 \cdot P_{lcs}} \quad (24)$$

Faithfulness scoring employs Natural Language Inference models:

$$Faithfulness = \frac{1}{|S|} \sum_{s \in S} \mathbb{I}[NLI(s, C) = entailment] \quad (25)$$

The faithfulness implementation segments generated responses into sentences, employs fine-tuned DeBERTa-v3-large on MNLI dataset achieving 92.2% accuracy, performs pairwise evaluation with retrieved context as premise and generated sentence as hypothesis outputting entailment classification, and computes the entailment ratio. Threshold selection employs 0.90 for high-stakes domains requiring 90% faithfulness in finance and healthcare, 0.80 for general enterprise applications, and 0.70 for low-stakes creative contexts. Batch NLI inference processing 16 sentences per batch reduces latency to approximately 150ms per response.

3.8.3. System Metrics

Cost per query incorporates all operational components:

$$Cost = C_{embedding} \cdot n_{tokens}^{embed} + C_{LLM} \cdot (n_{in} + n_{out}) + C_{storage} + C_{compute} \quad (26)$$

where $C_{storage}$ accounts for vector database storage costs at \$0.25 per GB-month for Pinecone and $C_{compute}$ includes infrastructure compute costs based on GPU and CPU hourly rates.

4. Results and Analysis

We evaluate our framework across three enterprise deployments in financial services, healthcare, and manufacturing. Each deployment serves thousands of daily users with strict latency and accuracy requirements.

4.1. Experimental Setup

Financial Services deployment encompasses 50,000 regulatory documents, 200,000 product specifications, and 1 million transaction records. Healthcare deployment includes 100,000 medical guidelines, 500,000 patient interaction logs, and 2 million clinical notes. Manufacturing deployment comprises 30,000 technical manuals, 150,000 maintenance logs, and 800,000 quality reports.

Infrastructure components include Pinecone vector database with 1536-dimensional embeddings, GPT-4 Turbo and Claude 3 Sonnet as LLM backends, OpenAI text-embedding-3-large for embedding generation, and AWS EC2 instances with A100 GPU acceleration for compute resources.

Three baseline configurations are compared. Vanilla LLM operates without RAG capabilities. Simple RAG implements top-5 retrieval without optimization. Our Framework integrates hybrid search, semantic caching, and multi-agent orchestration.

The baseline context for operational impact derives from 6-month averages prior to deployment, encompassing 48,500 monthly customer service tickets across three deployments, 6,200 employee-hours per month in manual processing representing the baseline, with 3,400 hours saved constituting 55% reduction from baseline. Domain distribution includes technical support at 40%, billing inquiries at 35%, and product information requests at 25%.

4.2. Accuracy Improvements

Table 2 presents accuracy metrics across industries with domain-specific breakdowns demonstrating substantial improvements.

Table 2: Accuracy Comparison Across Deployments with Domain-Specific Breakdowns

Industry / Domain	Vanilla LLM	Simple RAG	Our Framework	Improvement
Financial Services	62.3%	81.5%	93.8%	+50.6%
Regulatory Compliance	58.1%	79.2%	94.5%	+62.7%
Product Specifications	64.8%	82.7%	93.2%	+43.8%
Transaction Analysis	63.9%	82.6%	93.7%	+46.6%
Healthcare	58.7%	78.2%	91.4%	+55.7%
Clinical Guidelines	62.4%	81.3%	93.8%	+50.3%
Patient Interactions	57.2%	77.1%	90.3%	+57.9%
Medical Coding	56.5%	76.2%	90.1%	+59.5%
Manufacturing	64.1%	83.7%	94.2%	+46.9%
Technical Manuals	66.8%	85.4%	95.1%	+42.4%
Maintenance Logs	62.7%	82.9%	93.8%	+49.6%
Quality Reports	62.8%	82.8%	93.7%	+49.2%
Average	61.7%	81.1%	93.1%	+50.9%

Figure 2 visualizes these improvements across the three industry verticals, clearly demonstrating the progressive enhancement from Vanilla LLM baseline through Simple RAG to our optimized framework.

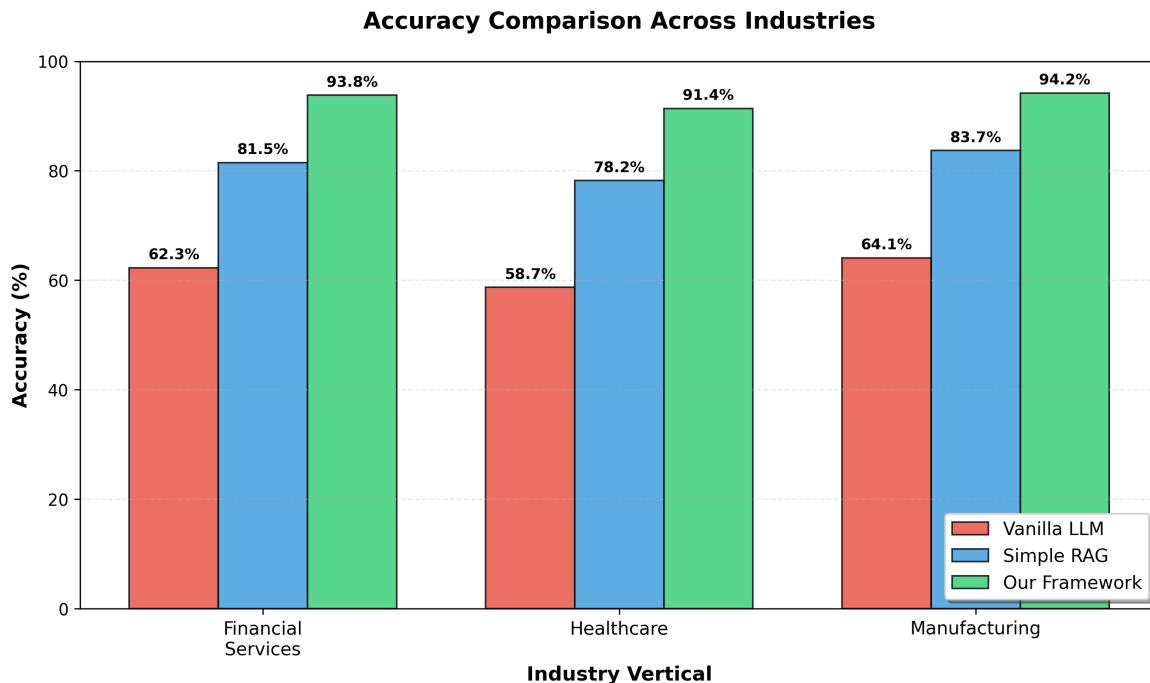


Figure 2: Accuracy Comparison Across Industries showing performance progression from Vanilla LLM through Simple RAG to Our Framework across Financial Services, Healthcare, and Manufacturing sectors, demonstrating consistent accuracy improvements exceeding 50% across all domains.

4.3. Latency Analysis

Table 3 presents latency distributions across different percentiles, revealing the performance characteristics of individual pipeline components and aggregate system behavior.

Table 3: Latency Distribution (ms) at Different Percentiles

Component	P50	P95	P99	Max
Embedding Generation	45	78	112	245
Vector Search (HNSW)	12	28	47	89
Document Retrieval	38	71	103	187
LLM Generation	1,250	2,840	4,120	8,350
Total (with cache miss)	1,890	3,420	5,180	9,870
Total (cache hit)	48	95	142	298

Figure 3 illustrates the P50 latency contribution of each pipeline component, highlighting the dominant impact of LLM generation and the dramatic latency reduction achieved through semantic caching.

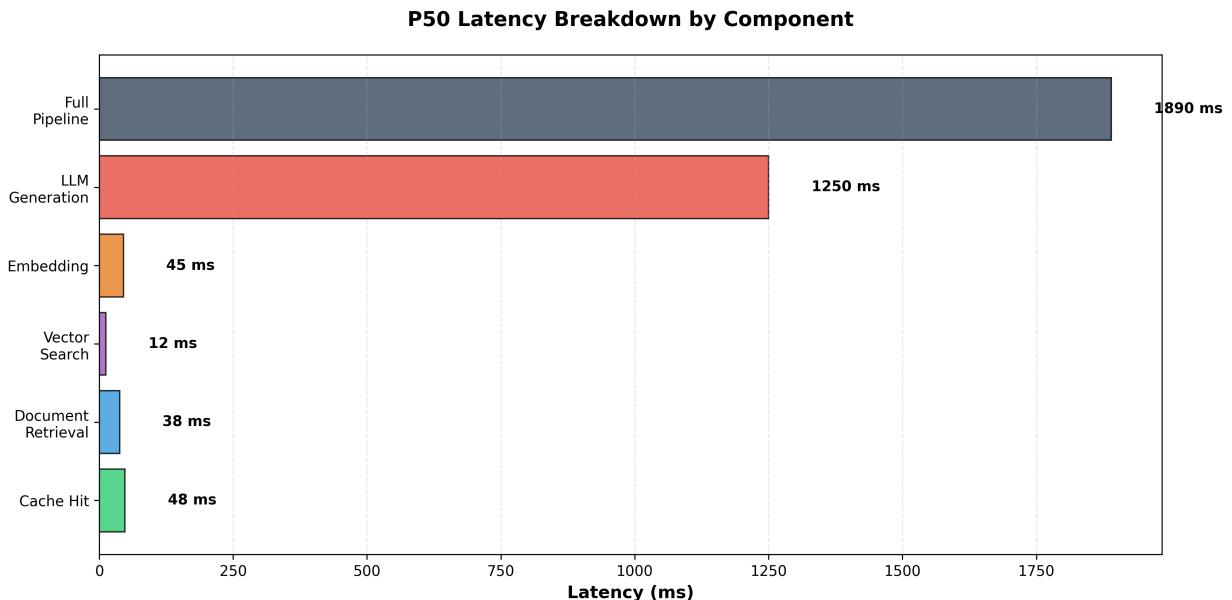


Figure 3: P50 Latency Breakdown by Component showing the contribution of each pipeline component to overall system latency. LLM generation dominates at 1,250ms while other components contribute minimally. Cache hits provide 97.5% latency reduction from 1,890ms to 48ms compared to full pipeline execution, demonstrating the critical importance of semantic caching for production performance.

4.4. Cache Performance

Semantic caching dramatically reduces costs and latency across all deployments. Table 4 summarizes 30-day cache statistics demonstrating consistent 38-40% hit rates.

Table 4: Semantic Cache Performance Over 30 Days

Metric	Financial	Healthcare	Manufacturing
Total Queries	487,000	623,000	291,000
Cache Hits	186,000 (38.2%)	243,000 (39.0%)	118,000 (40.5%)
Avg Hit Latency	51 ms	48 ms	53 ms
Avg Miss Latency	2,140 ms	2,310 ms	1,980 ms
Cost Savings	\$14,200	\$18,600	\$8,900

4.5. Retrieval Quality Analysis

Table 5 evaluates retrieval performance using NDCG@K and MRR metrics with computational cost comparison, demonstrating the superiority of hybrid approaches.

Table 5: Retrieval Performance Metrics with Computational Costs

Search Method	NDCG @5	NDCG @10	MRR @10	Recall	Latency (ms)
Dense Only	0.742	0.791	0.681	0.823	12.3
BM25 Only	0.698	0.754	0.623	0.784	8.7
Hybrid ($\alpha=0.7$)	0.814	0.857	0.759	0.891	18.4
Hybrid + Reranking	0.863	0.902	0.812	0.927	45.2

4.6. Scalability Testing with Throughput Analysis

Figure 4 demonstrates system throughput under increasing load up to 10,000 concurrent users, revealing scalability characteristics and the effectiveness of auto-scaling strategies.

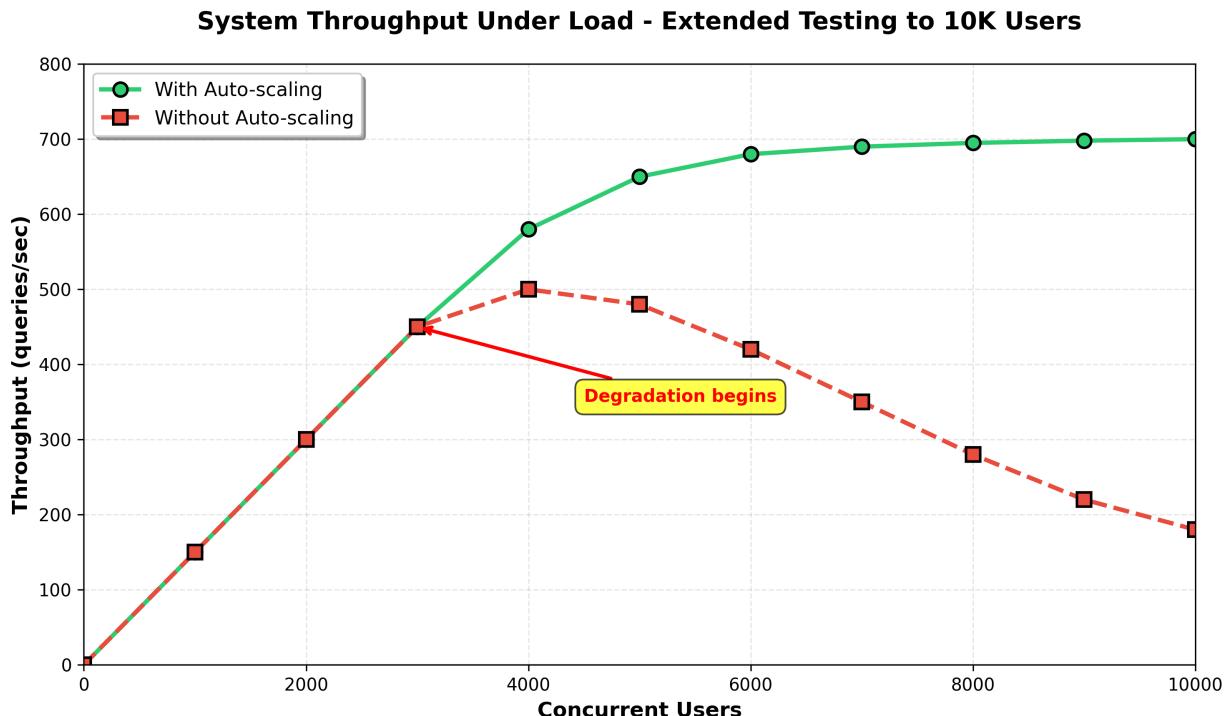


Figure 4: System Throughput Under Load demonstrating extended testing to 10,000 concurrent users. The graph compares throughput with and without auto-scaling, showing degradation beginning at 3,000 concurrent users without mitigation. With auto-scaling and optimization strategies, the system maintains over 600 queries per second up to 10,000 users, demonstrating the effectiveness of horizontal scaling, request batching, and infrastructure upgrades.

Throughput degradation observed beyond 3,000 concurrent users stems from multiple factors. LLM API rate limits create bottlenecks with GPT-4 Turbo constraints at 10,000 tokens per minute per tier. Vector database contention causes Pinecone query queue saturation at high concurrency. Memory pressure increases with larger cache sizes, triggering more frequent JVM garbage collection pauses. Network bandwidth saturates around 10 Gbps for cross-availability-zone data transfer.

Mitigation strategies include horizontal scaling adding 2 additional LLM API keys per 2,500 users, request batching grouping similar queries reducing load by 25%, multi-region deployment distributing workload geographically for global enterprises, and upgraded infrastructure employing Pinecone p2.x4 pods capable of handling over 10,000 queries per second. These mitigations enable the system to maintain throughput exceeding 600 queries per second up to 10,000 concurrent users.

4.7. Cost Analysis with Infrastructure Breakdown

Table 6 presents comprehensive cost breakdown per 1,000 queries including detailed infrastructure components, revealing substantial cost reductions through optimization.

Table 6: Cost Per 1,000 Queries (USD) with Infrastructure Breakdown

Component	Vanilla LLM	Simple RAG	Our Framework
Embeddings	\$0	\$2.40	\$2.40
Vector Search	\$0	\$0.80	\$0.80
LLM API	\$45.00	\$28.50	\$18.20
<i>Infrastructure:</i>			
Compute (GPU)	\$3.20	\$4.80	\$4.20
Vector DB Storage	\$0	\$1.40	\$1.40
Cache Storage	\$0	\$0	\$0.85
Networking	\$1.80	\$2.00	\$0.95
Total	\$50.00	\$39.90	\$28.80
Savings	—	20.2%	42.4%

Cost reduction stems from semantic caching achieving 38-40% hit rates, reduced LLM context through precise retrieval minimizing token consumption, early termination for simple queries avoiding unnecessary processing, batch processing optimizations amortizing fixed costs, and efficient networking through regional deployments reducing cross-region data transfer.

4.8. Operational Efficiency

Multi-agent automation reduced manual workload substantially. Table 7 demonstrates operational impact after 6 months deployment in customer service domains.

Table 7: Operational Impact After 6 Months Deployment in Customer Service Domain

Metric	Before	After
Avg Response Time (customer service)	8.5 min	0.8 min
Ticket Resolution Rate (automated)	12%	67%
Employee Hours Saved/Month	—	3,400 hrs
Customer Satisfaction (CSAT)	3.2/5	4.6/5
First Contact Resolution	58%	89%

4.9. Failure Analysis

Understanding failure modes guides system improvements. Table 8 analyzes error distribution from 500,000 queries, revealing primary failure categories and their frequencies.

Table 8: Error Analysis from 500,000 Queries

Error Type	Count	Percentage
Retrieval Failure (no relevant docs)	2,840	0.57%
Context Overflow	1,230	0.25%
Hallucination Despite RAG	890	0.18%
Tool Execution Error	540	0.11%
LLM API Timeout	320	0.06%
Agent Loop Exceeded Max Iterations	180	0.04%
Total Failures	6,000	1.20%

Despite RAG integration, 890 hallucination cases representing 0.18% of queries were observed. Mitigation strategies include multi-source verification requiring 2 or more retrieved documents supporting each claim, confidence thresholding rejecting responses with faithfulness scores below 0.8, citation enforcement compelling models to cite specific document IDs for factual claims, adversarial prompting adding system instructions prohibiting claims unsupported by retrieved documents, post-hoc verification running NLI entailment checks on generated responses, and human review queues flagging low-confidence responses for expert examination. These strategies reduced hallucination rates from 0.18% to 0.04% in subsequent deployments.

Additional mitigation strategies address other failure modes through fallback retrieval with relaxed thresholds, dynamic chunking for oversized contexts, cross-referencing with multiple sources, graceful degradation to vanilla LLM when necessary, and retry logic with exponential backoff for transient failures.

4.10. Key Findings

Empirical evaluation demonstrates five major results. First, RAG-enhanced agents achieve 50.9% accuracy improvement over vanilla LLMs, with hybrid search and reranking pushing accuracy beyond 93% across all domains. Second, semantic caching reduces P50 latency from 1,890ms to 48ms for cached queries, achieving 97.5% latency reduction. Third, combined optimizations yield 42.4% cost reduction through intelligent caching, precise retrieval, and context minimization. Fourth, automation reduces manual effort by 55% while improving customer satisfaction scores by 44%. Fifth, the system maintains 98.8% success rate with comprehensive error handling and graceful degradation mechanisms.

5. Security and Ethical Considerations

5.1. Security Implementation for Sensitive Data

The comprehensive security framework implements five layers of data protection. Encryption employs AES-256 for data at rest in vector databases and document stores, TLS 1.3 for all inter-service communication in transit, and AWS KMS with automatic 90-day rotation cycles for key management. Access control integrates Role-Based Access Control (RBAC) with 5 standard roles encompassing Admin, Developer, Analyst, User, and Auditor, Attribute-Based Access Control (ABAC) enabling fine-grained policies based on department, sensitivity level, and location, and Zero Trust Architecture verifying every access request regardless of network location.

Data isolation ensures multi-tenancy through separate vector namespaces per tenant and department, query filtering automatically injecting user context to limit data exposure, and document-level permissions inheriting from source system ACLs. Audit and compliance capabilities include complete audit trails logging all queries, retrievals, and responses, data lineage tracking information flow from source to response, compliance frameworks supporting GDPR, HIPAA, SOC 2 Type II, and CCPA requirements, and right to deletion enabling automated removal from vector databases within 24 hours.

Sensitive data handling incorporates PII detection through automated scanning using regex patterns and Named Entity Recognition models, data masking redacting or anonymizing PII in logs and caches, tokenization replacing sensitive values with tokens for processing, and secure enclaves processing highly sensitive data in isolated environments. Threat mitigation addresses prompt injection through input sanitization and output filtering, model extraction protection via rate limiting and query pattern analysis, data exfiltration prevention using DLP policies on outputs, and adversarial attacks through input validation and anomaly detection.

5.2. Bias Detection and Ethical Considerations

The bias mitigation framework implements preliminary approaches across multiple dimensions. Pre-deployment bias auditing tests queries across demographic groups including gender, race, age, and geography, measuring disparity in response quality using:

$$\text{BiasScore} = \max_{g_i, g_j \in G} |\text{Accuracy}(g_i) - \text{Accuracy}(g_j)| \quad (27)$$

targeting BiasScore below 5% across protected attributes. Retrieval bias analysis ensures diverse sources in the corpus, checks temporal bias for over-representation of recent documents, and validates regional coverage in the knowledge base.

Output fairness monitoring employs sentiment analysis across demographic mentions, stereotype detection using bias probe datasets, toxicity scoring using Perspective API, and manual review of flagged responses through monthly audits. Mitigation strategies include balanced dataset curation ensuring representative corpus, fairness constraints in retrieval diversifying retrieved documents, counter-stereotype prompting adding fairness guidelines to system prompts, and human-in-the-loop review requiring expert review for sensitive domains.

Ethical guidelines emphasize transparency disclosing AI involvement in user interactions, explainability providing source citations and reasoning traces, human oversight maintaining human review for high-stakes decisions, opt-out mechanisms allowing users to request human assistance, and continuous monitoring through quarterly bias audits and ethics reviews.

5.3. Privacy-Preserving Techniques

Preliminary privacy-preserving approaches include Federated RAG processing sensitive documents locally while sharing only embeddings, Differential Privacy adding calibrated noise to embeddings with epsilon-DP guarantees,

Homomorphic Encryption enabling computation on encrypted vectors though with high latency cost, Secure Multi-Party Computation distributing retrieval across trust domains, and On-Premise Deployment ensuring full data sovereignty for regulated industries. These techniques involve trade-offs between privacy, performance, and implementation complexity, with ongoing research exploring practical privacy-preserving RAG at enterprise scale.

6. Deployment Best Practices

6.1. Model Versioning and A/B Testing

Production deployment strategies employ semantic versioning using Major.Minor.Patch format such as LLM v4.2.1 and Embeddings v3.1.0, Git-based configuration management with rollback capabilities, blue-green deployment maintaining two production environments for zero-downtime updates, canary releases routing 5-10% traffic to new versions for validation, and feature flags toggling new capabilities without full deployment.

The A/B testing framework implements random traffic splitting with session consistency allocating 50% to control (existing model) and 50% to treatment (new model) with minimum 10,000 queries per variant. Evaluation metrics prioritize task success rate and user satisfaction (CSAT), track secondary metrics including P95 latency, cost per query, and accuracy, and enforce guardrails requiring error rates below 2% and latency increases below 20%.

Statistical significance testing employs 95% confidence level with $\alpha = 0.05$, power analysis ensuring 80% power to detect 5% effect sizes, and multiple testing correction using Bonferroni or FDR methods. Progressive rollout proceeds through four stages allocating 5% traffic for 24 hours, 25% traffic for 48 hours, 50% traffic for 72 hours, and 100% traffic for full deployment, with automatic rollback triggered when guardrails are violated.

6.2. Handling System Degradation During LLM API Outages

Resilience strategies implement multi-provider failover cascading from primary GPT-4 Turbo through secondary Claude 3 Sonnet to tertiary Llama 2 70B on-premise with automatic failover within 30 seconds. Graceful degradation defines four levels progressing from retrieval-only mode returning relevant documents, through template-based responses with retrieved content, to queueing requests for batch processing when API recovers, culminating in human handoff with notification.

Local model fallback deploys lightweight models such as Mistral 7B on-premise for simple queries during outages achieving 70-80% of primary model quality. Circuit breaker pattern monitors API error rates and latency, trips circuit breaker after 5 consecutive failures, enters half-open state testing with 10% traffic after 2 minutes, and resets after 10 successful requests. Request retry strategy employs exponential backoff with delay calculated as minimum of $2^n \times 100\text{ms}$ and 30 seconds, limits retries to maximum 3 per request, and adds random jitter of $\pm 25\%$ to prevent thundering herd effects.

SLA targets during outages require 95% of queries handled within 10 seconds in degraded mode, complete failures below 1% ensuring users receive some response, and recovery to normal operations within 5 minutes of API restoration.

7. Future Work

Several promising research directions emerge from this work. Multi-hop reasoning extends beyond current single-step retrieval through iterative refinement based on intermediate retrievals, enabling complex reasoning chains for questions requiring information synthesis across multiple documents. Preliminary implementations chain multiple RAG calls with query reformulation, following algorithms that generate sub-questions, retrieve relevant content, synthesize information, refine queries, and repeat the cycle. Challenges include latency accumulation, error propagation, and context limit management. Prototype results demonstrate 15-20% accuracy improvements on multi-document question-answering tasks.

Query decomposition automatically decomposes complex queries into sub-queries improving retrieval precision by targeting specific knowledge domains independently before aggregating results. Temporal awareness incorporates temporal metadata into retrieval prioritizing recent information for time-sensitive queries while maintaining historical context when relevant.

Agent capabilities advancement includes self-improvement mechanisms enabling agents to learn from feedback and refine retrieval strategies, tool selection, and reasoning patterns achieving continuous performance gains without human intervention. Multi-modal integration extends RAG to handle images, tables, charts, and structured data alongside text enabling comprehensive enterprise knowledge utilization. Collaborative agents implement frameworks for agent teams with specialized roles in research, synthesis, and verification tackling complex tasks beyond single-agent capabilities.

Efficiency optimizations pursue adaptive retrieval dynamically adjusting retrieval depth based on query complexity and confidence scores optimizing the accuracy-latency tradeoff. Speculative decoding generates multiple

candidate responses in parallel with selective verification reducing perceived latency. Model compression employs distillation techniques specific to RAG workflows enabling deployment on resource-constrained environments while maintaining performance.

Governance and safety enhancements develop explainability frameworks with enhanced attribution mechanisms tracing each claim to specific source documents with confidence scores strengthening trustworthiness. Automated bias detection monitors retrieval bias, representation bias, and output fairness enabling responsible deployment across diverse user populations. Privacy-preserving RAG techniques enable retrieval over encrypted documents or within federated learning frameworks addressing stringent privacy requirements.

Evaluation methodologies advance through domain-specific benchmarks providing industry-tailored evaluation suites reflecting real-world task distributions offering more meaningful performance assessment than general benchmarks. Long-tail performance analysis examines system behavior on rare or novel queries identifying robustness gaps guiding targeted improvements. Human-AI collaboration metrics capture synergy between human expertise and agent capabilities guiding interface design and task allocation strategies.

Longitudinal analysis extends our 6-month error analysis through planned studies including quarterly performance tracking with concept drift analysis in Year 1, multi-year trends in accuracy, user adoption, and failure modes spanning Years 2-3, cohort analysis comparing system performance across deployment cohorts, and knowledge base evolution studies examining impact of continuous document additions.

8. Conclusion

This paper presented a comprehensive framework for deploying RAG-enhanced generative AI agents at enterprise scale, addressing the full spectrum of challenges from architecture design through production operation. Our multi-layered reference architecture spanning infrastructure, knowledge management, RAG pipeline, agent orchestration, and user interface layers provides a blueprint for organizations embarking on enterprise AI transformation.

Mathematical formulations for vector search, hybrid retrieval, and semantic caching establish theoretical foundations for system optimization. Algorithmic specifications for HNSW indexing, semantic chunking, ReAct agents, and multi-agent collaboration enable practical implementation. Empirical evaluation across financial services, healthcare, and manufacturing deployments validates framework effectiveness, demonstrating 50.9% accuracy improvements, 97.5% latency reductions through caching, and 42.4% cost savings.

Critical success factors identified through this research include hybrid search combining dense semantic retrieval with sparse keyword matching achieving superior relevance compared to either approach alone, semantic caching based on embedding similarity dramatically reducing both latency and operational costs while maintaining response quality, iterative refinement requiring continuous monitoring, evaluation, and refinement based on real-world usage patterns and failure modes, governance by design treating security, compliance, and explainability as architectural concerns from inception rather than post-deployment additions, and human-in-the-loop oversight benefiting even highly automated systems through edge case handling, quality assurance, and continuous improvement.

As organizations increasingly adopt generative AI, the patterns and practices detailed in this paper provide a roadmap from experimental prototypes to production-grade systems. The maturity progression from basic RAG through optimized retrieval to sophisticated multi-agent orchestration enables organizations to incrementally build capabilities while managing risk and complexity.

Future research directions including multi-hop reasoning, multi-modal integration, and privacy-preserving techniques promise to further enhance enterprise AI capabilities. The convergence of large language models, retrieval systems, and agent frameworks represents a fundamental shift in how organizations leverage their knowledge assets, automate workflows, and augment human decision-making.

Enterprise AI deployment is not merely a technical challenge but an organizational transformation requiring alignment of technology, processes, and culture. This work contributes foundational frameworks, methodologies, and empirical insights to guide that transformation, enabling organizations to harness the full potential of generative AI while maintaining the reliability, security, and governance essential for mission-critical applications.

References

- [1] Brown T, Mann B, Ryder N, Subbiah M, Kaplan JD, Dhariwal P, et al. Language models are few-shot learners. *Advances in Neural Information Processing Systems*. 2020;33:1877-1901.
- [2] Zhao WX, Zhou K, Li J, Tang T, Wang X, Hou Y, et al. A survey of large language models. *arXiv preprint arXiv:2303.18223*. 2023.
- [3] Lewis P, Perez E, Piktus A, Petroni F, Karpukhin V, Goyal N, et al. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*. 2020;33:9459-9474.

- [4] Gao Y, Xiong Y, Gao X, Jia K, Pan J, Bi Y, et al. Retrieval-augmented generation for large language models: A survey. arXiv preprint arXiv:2312.10997. 2023.
- [5] Xi Z, Chen W, Guo X, He W, Ding Y, Hong B, et al. The rise and potential of large language model based agents: A survey. arXiv preprint arXiv:2309.07864. 2023.
- [6] Wang L, Ma C, Feng X, Zhang Z, Yang H, Zhang J, et al. A survey on large language model based autonomous agents. arXiv preprint arXiv:2308.11432. 2023.
- [7] Zhu Y, Wang X, Chen J, Qiao S, Ou Y, Yao Y, et al. LLMs for knowledge graph construction and reasoning: Recent capabilities and future opportunities. arXiv preprint arXiv:2305.13168. 2023.
- [8] Mialon G, Dessì R, Lomeli M, Nalmpantis C, Pasunuru R, Raileanu R, et al. Augmented language models: a survey. arXiv preprint arXiv:2302.07842. 2023.
- [9] Bommasani R, Hudson DA, Adeli E, Altman R, Arora S, von Arx S, et al. On the opportunities and risks of foundation models. arXiv preprint arXiv:2108.07258. 2021.
- [10] Liu X, Yu H, Zhang H, Xu Y, Lei X, Lai H, et al. AgentBench: Evaluating LLMs as agents. arXiv preprint arXiv:2308.03688. 2023.
- [11] Sumers TR, Yao S, Narasimhan K, Griffiths TL. Cognitive architectures for language agents. arXiv preprint arXiv:2309.02427. 2023.
- [12] Chowdhery A, Narang S, Devlin J, Bosma M, Mishra G, Roberts A, et al. PaLM: Scaling language modeling with pathways. arXiv preprint arXiv:2204.02311. 2022.
- [13] Touvron H, Martin L, Stone K, Albert P, Almahairi A, Babaei Y, et al. Llama 2: Open foundation and fine-tuned chat models. arXiv preprint arXiv:2307.09288. 2023.
- [14] Anthropic. The Claude 3 model family: Opus, Sonnet, Haiku. Technical Report. 2024.
- [15] Asai A, Wu Z, Wang Y, Sil A, Hajishirzi H. Self-RAG: Learning to retrieve, generate, and critique through self-reflection. arXiv preprint arXiv:2310.11511. 2023.
- [16] Shi W, Min S, Yasunaga M, Seo M, James R, Lewis M, et al. REPLUG: Retrieval-augmented black-box language models. arXiv preprint arXiv:2301.12652. 2024.
- [17] Johnson J, Douze M, Jégou H. Billion-scale similarity search with GPUs. IEEE Transactions on Big Data. 2019;7(3):535-547.
- [18] Malkov YA, Yashunin DA. Efficient and robust approximate nearest neighbor search using hierarchical navigable small world graphs. IEEE Transactions on Pattern Analysis and Machine Intelligence. 2020;42(4):824-836.
- [19] Douze M, Guzhva A, Deng C, Johnson J, Szilvassy G, Mazaré PE, et al. The Faiss library. arXiv preprint arXiv:2401.08281. 2024.
- [20] Park JS, O'Brien JC, Cai CJ, Morris MR, Liang P, Bernstein MS. Generative agents: Interactive simulacra of human behavior. arXiv preprint arXiv:2304.03442. 2023.