Published by :
https://www.ijert.org/
An International Peer-Reviewed Journal

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 10 , October - 2026

AI-Powered Virtual CFO and Compliance Automation Platform with Intelligent Document Vault

A Multi-Agent Architecture and Its Implementation as FinVault AI

Likhil Gowda K A
Dept. of CSE (Data Science)
AMC Engineering College
Bengaluru, Karnataka, India

Bharath N
Dept. of CSE (Data Science)
AMC Engineering College
Bengaluru, Karnataka, India

Madhan V
Dept. of CSE (Data Science)
AMC Engineering College
Bengaluru, Karnataka, India

Mohith Nagdeep
Dept. of CSE (Data Science)
AMC Engineering College
Bengaluru, Karnataka, India

Prof. Swagatika Samal
Assistant Professor
Dept. of CSE (Data Science)
AMC Engineering College
Bengaluru, Karnataka, India

Abstract—Financial, compliance, and document-management work in small businesses, accounting firms, and households remains largely manual and fragmented across spreadsheets, invoices, bank statements, messaging applications, and file folders. This paper presents the design and implementation of an AI-powered Virtual CFO and Compliance Automation Platform with an Intelligent Document Vault, a multi-agent system that brings financial analysis, compliance assistance, and personal document intelligence into one secure architecture. An AI CFO Agent and a Compliance Agent serve chartered accountants, accountants, and businesses through a web application, supporting profit and loss, cash-flow, and budget-versus-actual analysis, forecasting, deadline tracking, and missing-document detection, while a DocVault Agent lets individual users store and query personal documents through WhatsApp using optical character recognition and retrieval-augmented generation. Large language models are used for reasoning and explanation, whereas deterministic backend tools perform every financial calculation. The agents are orchestrated with LangGraph, and security is enforced through role-based access control, tenant isolation, audit logging, and human approval of sensitive outputs. The web portal has been implemented and deployed as FinVault AI, and its dashboard, AI CFO workspace, and compliance agent are described. A plan for the quantitative evaluation of forecasting, extraction, retrieval, and security is also given.

Keywords—AI CFO; compliance automation; document vault; agentic AI; retrieval-augmented generation; optical character recognition; financial forecasting; LangGraph; human-in-the-loop

I. INTRODUCTION

Small businesses, accounting firms, and individual users still perform a large share of their financial and compliance work manually. Invoices, bank statements, bills, and spreadsheets are spread across separate tools, which makes it difficult to obtain a complete and timely view of a business's position. Compliance work adds the continuing task of monitoring deadlines, required documents, and client-specific obligations. Individuals face a similar problem with insurance policies, warranties, certificates, and vehicle records that are scattered across phones, cloud drives, and messaging applications.

Large Language Models (LLMs) built on the Transformer architecture [1], [2] now make natural-language interaction with such information practical. However, an LLM used alone is unsuitable for financial work: it can produce fluent but unsupported statements and it is unreliable at exact arithmetic. Several complementary techniques address these weaknesses. Retrieval-Augmented Generation (RAG) grounds answers in a user's own records [3], Optical Character Recognition (OCR) engines such as Tesseract make scanned documents searchable [4], agent patterns such as ReAct allow a model to call external tools instead of computing results itself [5], and ensemble learners such as Random Forest [6] and XGBoost [7] provide established methods for tabular forecasting.

This paper presents the design and implementation of an AI-powered Virtual CFO and Compliance Automation Platform with an Intelligent Document Vault. The platform is a multi-agent system with three specialised agents: an AI CFO Agent and a Compliance Agent for chartered accountants (CAs), accountants, and businesses using a web application, and a DocVault Agent that gives individual users personal document intelligence through WhatsApp. Unlike a general-purpose chatbot, the system executes structured workflows covering calculations, document processing, report generation, and task tracking, and it holds sensitive outputs for human approval.

The main contributions of this work are:

- a three-agent architecture that serves business and personal users on shared infrastructure;
- a design rule that confines the LLM to interpretation and explanation, while deterministic backend tools perform all financial calculations;
- a seven-layer secure architecture with role-based access control, tenant isolation, audit logging, and human-in-the-loop approval by the CA;
- an implemented and deployed web portal, FinVault AI, with a dashboard, an AI CFO workspace with forecasting, and a compliance agent; and

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 10 , October - 2026

- a requirements specification and an evaluation plan covering forecasting, OCR, retrieval, and security.

Quantitative performance results have not yet been collected, so none are claimed here; the evaluation plan in Section IX defines how they will be obtained. The rest of the paper is organised as follows. Section II reviews related work and Section III states the problem. Sections IV and V present the proposed system and its architecture, and Section VI describes its workflows. Section VII describes the implemented web portal. Section VIII specifies the requirements and security controls, Section IX gives the methodology and evaluation plan, and Section X concludes with limitations and future work.

II. RELATED WORK

A. Machine Learning for Forecasting

Random Forest averages many decision trees trained on random subsets of the data and features, which improves robustness over a single tree [6]. XGBoost is a scalable gradient-boosting system that is widely used for tabular prediction [7]. Either method can support revenue, expense, and cash-flow forecasting when enough reliable historical data exists. Forecasts need careful evaluation with error measures such as Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and Mean Absolute Percentage Error (MAPE), and they should be presented as estimates based on past patterns, not as guaranteed outcomes.

B. Language Models and Agents

The Transformer uses attention to model relationships within a sequence [1], and large pretrained language models were later shown to perform varied language tasks from a few examples given in the prompt [2]. ReAct interleaves reasoning steps with tool actions so that a model can interact with external tools [5]. In the proposed platform an agent follows this idea: it interprets a request and routes it to a financial-calculation, document-search, or compliance function, and validated tools rather than the model produce the critical numbers. Generated text must still be checked against reliable sources and system rules.

C. Retrieval-Augmented Generation and OCR

A language model cannot know a user's current documents. RAG addresses this by retrieving relevant passages from an external store and supplying them during answer generation [3], which also makes answers more traceable to evidence. OCR converts scanned pages and images into machine-readable text, and Tesseract is an established open-source engine for this purpose [4]. Extraction accuracy depends on image quality, layout, language, and document condition, so extracted values should be validated before they drive reminders or reports.

D. Research Gap

Each technique contributes one capability: forecasting, language interaction, tool use, grounded retrieval, or document digitisation. None of them alone solves financial, compliance, and document-management work. The gap addressed here is their combination in one governed platform, in which every technique is confined to the task it suits and a human remains accountable for sensitive outputs.

III. PROBLEM STATEMENT AND EXISTING SYSTEM

A. Existing Practice

Small businesses and accounting firms currently handle financial and compliance work through spreadsheets, accounting applications, e-mail, messaging platforms, and manual records. Staff spend considerable time collecting documents, entering transactions, and preparing summaries and Management Information System (MIS) reports. Because information is scattered across invoices, bank statements, and other files, consolidated analysis is difficult, and manual tracking makes it hard to know which documents are available, which tasks are pending, and which deadlines are close. Conventional document storage has no intelligence: users must open files to find details such as an insurance expiry date or a warranty period.

B. Limitations

The main limitations of existing approaches are:

- repetitive manual work in data collection, entry, and report preparation;
- scattered data that prevents a consolidated view of performance;
- little or no predictive capability beyond historical reporting;
- compliance deadlines, documents, and client records tracked separately;
- no natural-language retrieval from stored documents; and
- no unified workflow that combines AI, machine learning, document processing, calculation, and human approval.

C. Problem Statement

There is a need for an intelligent system that unifies financial processing, compliance management, document understanding, forecasting, and retrieval, and that reduces repetitive work while keeping calculations reliable. The central requirement is to separate language-model reasoning from deterministic financial tools instead of relying on model-generated arithmetic. Table I contrasts existing practice with the proposed platform.

TABLE I. EXISTING PRACTICE VERSUS THE PROPOSED PLATFORM

Aspect	Existing practice	Proposed platform
Financial analysis	Manual entry in spreadsheets or basic accounting tools	Deterministic tools for P&L, cash flow, and budget variance
Forecasting	Little or none	Random Forest and XGBoost forecasts with stated limits
Compliance	Spreadsheets, e-mail, manual follow-up	Calendar, checklists, missing-document detection
Document retrieval	Manual folder search	OCR and RAG question answering over WhatsApp
Integration	Separate tools	One multi-agent platform on shared services
Oversight	Ad hoc	CA approval and audit logs

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 10 , October - 2026

IV. PROPOSED SYSTEM

The proposed platform is a multi-agent system with three specialised agents. They share common infrastructure but differ in responsibilities and user experience.

A. AI CFO Agent

The AI CFO Agent acts as a digital financial-operations assistant for CAs, accountants, and businesses. It receives structured financial data and documents, calls deterministic tools to compute Profit and Loss (P&L), cash flow, receivables and payables, and budget-versus-actual variance, and identifies anomalies and trends. A forecasting component trains Linear Regression, Random Forest [6], and XGBoost [7] models on historical revenue, expense, and cash data. The agent then explains the predictions and their limits in generated reports such as the monthly MIS report.

B. Compliance Agent

The Compliance Agent supports CAs and businesses with preparation and organisation rather than autonomous filing. It maintains a compliance calendar, tracks upcoming and overdue tasks, checks required documents against client-wise checklists, identifies missing items, and prepares status reports and reminders. Regulated filings, tax payments, and final decisions remain with qualified users.

C. DocVault Agent

The DocVault Agent is a separate personal-document module for normal users, accessed through WhatsApp. Users upload documents such as warranties, insurance policies, certificates, vehicle records, and bills. The agent extracts text with OCR, identifies metadata, splits the text into chunks, creates embeddings, and indexes them for semantic search. Questions such as “When does my insurance expire?” are answered through RAG [3] from the user’s own documents, and expiry reminders are sent where dates can be reliably identified.

D. Shared Infrastructure and Technology Stack

The agents are orchestrated with LangGraph [8], which expresses each workflow as a controlled multi-step graph with explicit tool permissions. Table II lists the technology stack.

TABLE II. TECHNOLOGY STACK

Component	Technology
Web front end	Next.js
API gateway	FastAPI [9]
Orchestration	LangGraph [8]
Structured data	PostgreSQL
Vector search	pgvector [10]
Object storage	S3 / Supabase
Cache	Redis
Reports and tools	Pandas, ReportLab, openpyxl
Forecasting models	Random Forest, XGBoost
Deployment	Docker; AWS / Railway / Render

V. SYSTEM ARCHITECTURE

The architecture, shown in Fig. 1, has seven layers. Side panels show external systems, cross-cutting concerns, data governance, and the human-in-the-loop approval step.

- **User layer.** A responsive web application for CAs and business users, and a WhatsApp chatbot for normal users.
- **API gateway and security layer.** A FastAPI gateway with JWT/OAuth 2.0 authentication, role-based authorisation, tenant isolation, and a WhatsApp webhook.
- **Agent and application layer.** The LangGraph orchestrator, the three agents, and approval and notification services.
- **Intelligence layer.** LLM reasoning and generation, the RAG engine, forecasting models, OCR and data extraction, and insights.
- **Tool layer.** Deterministic functions: financial tools, a compliance rules engine, PDF and Excel report generation, and reminder jobs.
- **Data layer.** PostgreSQL for structured records, pgvector for embeddings, object storage for documents and reports, audit logs and approvals, and a cache.
- **Infrastructure and deployment layer.** Docker containers, background workers, cloud hosting, a CI/CD pipeline, and monitoring.

Background workers handle long-running tasks such as OCR, indexing, forecasting, and report generation so that the interface stays responsive.

VI. SYSTEM DESIGN AND WORKFLOWS

A. Actors and Use Cases

Three actors use the platform. The CA or accountant manages clients, imports transactions, runs P&L, cash-flow, and MIS reports, views forecasts, reviews and approves outputs, tracks compliance tasks, and checks missing documents. The business user views financial summaries, uploads bills and invoices, compares budget with actual figures, downloads PDF or Excel reports, and receives due-date reminders. The normal user uploads documents on WhatsApp, asks questions about them, checks expiry and renewal dates, searches stored documents, and receives reminders. Login, role-based access control, human approval, notifications, OCR and indexing, the RAG assistant, and the expiry-reminder scheduler are shared functions used by all three.

B. Data Flow

Data enter the platform as uploads, transactions, WhatsApp media, and queries. Input processing performs the authentication and tenant check, file import, field validation, and media download. OCR and extraction then obtain text, tables, and metadata, chunk and embed the content, and classify documents. The results flow to the DocVault engine (RAG answering, expiry detection, reminders, and personal search), the AI CFO engine (P&L, cash flow, receivables and payables, budget variance, and expense analysis), and the Compliance engine (calendar and rules, required-document checks, missing-item detection, and client-wise status). The

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 10 , October - 2026

forecasting model and the LLM insight layer then produce forecasts and explanations, and the LLM performs no free-form arithmetic. The report generator, the approval and audit process, and the user interfaces deliver MIS reports, Excel exports, dashboards, WhatsApp replies, and alerts, while the operational database, vector store, document storage, and audit log are accessed throughout.

C. Request Lifecycle

Fig. 2 shows how a request moves through the platform. First, the user submits an upload, report request, or question through the web application or WhatsApp. The API gateway forwards it to the authentication service, which checks the token, role, and tenant; an invalid request returns an error. A valid request reaches the agent orchestrator, which plans the workflow and selects the agent and tool.

For a financial or compliance request, the selected tool retrieves transactions and rules from the database and returns a computed result. For a document question, the retrieval tool fetches relevant chunks from the vector store, using OCR-extracted text for scanned files [3], [4]. The LLM then drafts an explanation of the verified results, and the result is saved with an audit entry. For sensitive financial or compliance outputs, the draft is shown to the CA, who approves or rejects it; only the approved version is finalised and stored with its audit trail. Every step is recorded in the audit and agent-run tables.

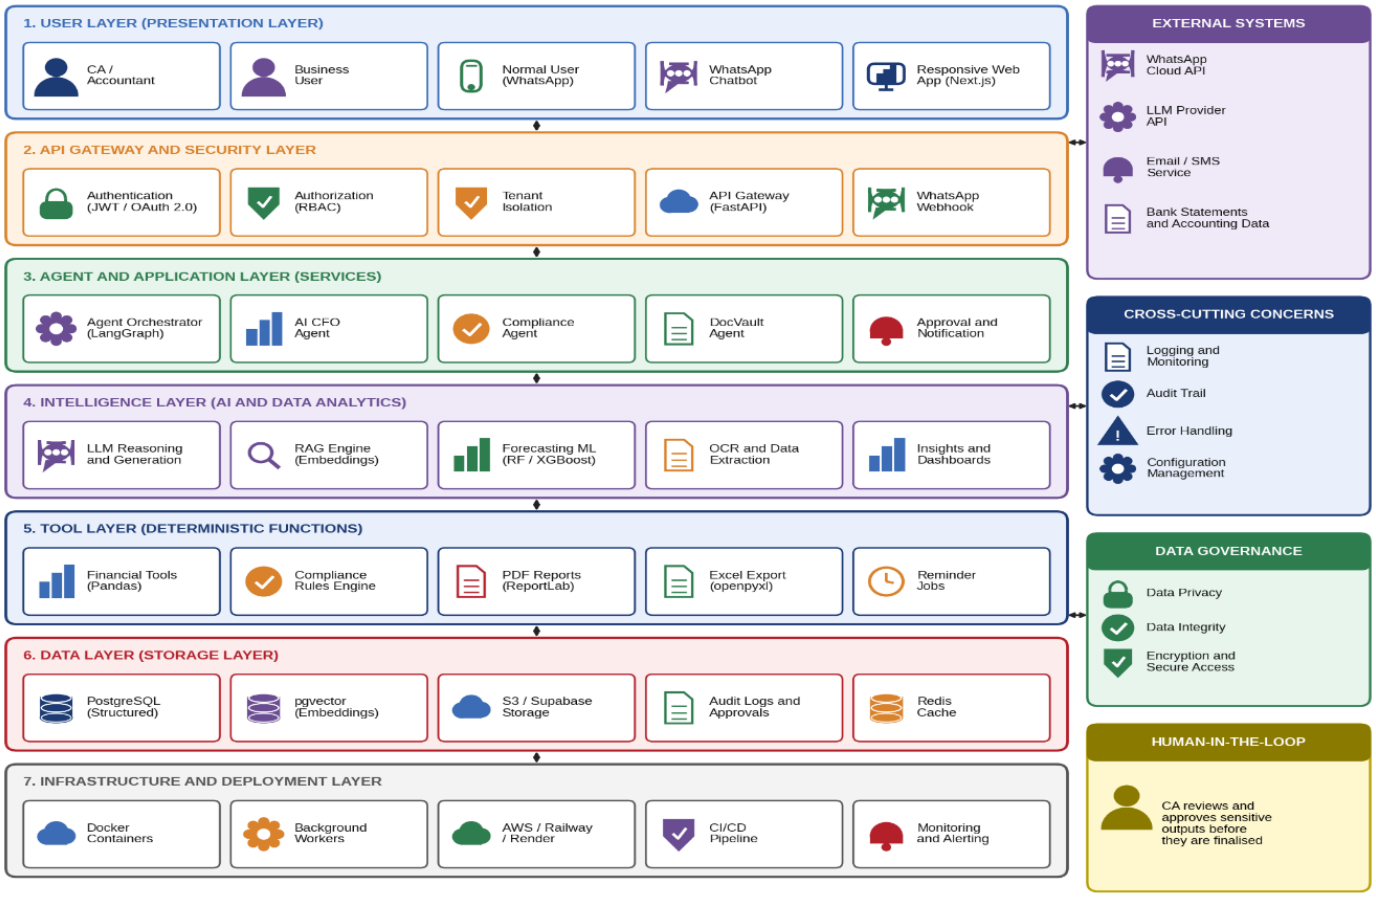


Fig. 1. Seven-layer architecture of the proposed platform, with external systems, cross-cutting concerns, data governance, and human-in-the-loop approval.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 10 , October - 2026

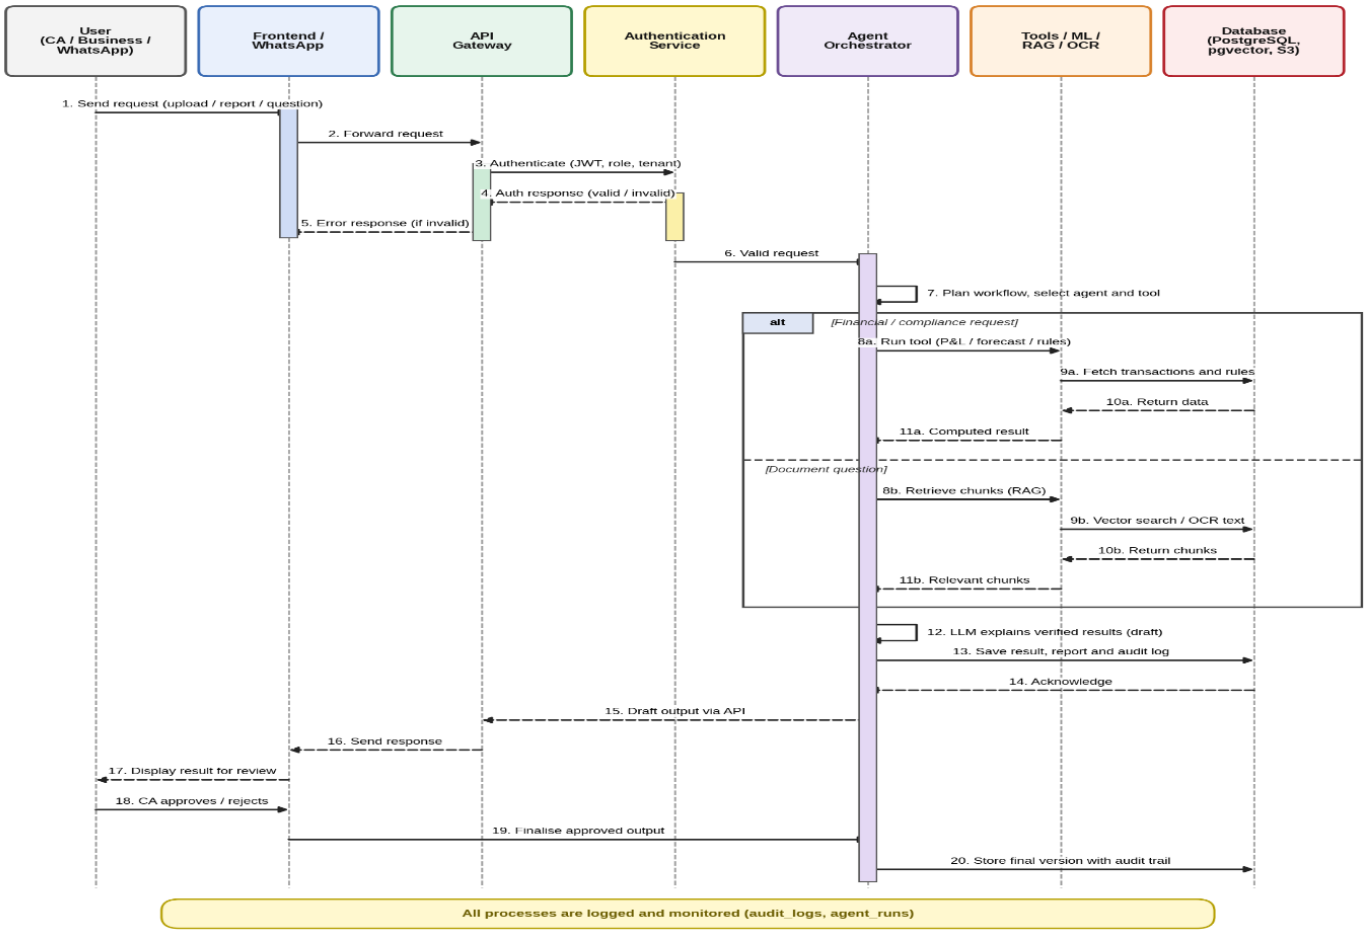


Fig. 2. Sequence diagram of the request lifecycle, from the user's request to CA approval and audit storage.

VII. IMPLEMENTATION

The design has been implemented and deployed as a web application named FinVault AI (Finance and Compliance) for the platform administrator, CAs, and business users. Figs. 3 to 5 show its main screens for an example client organisation, ABC Technologies Pvt Ltd. They are screenshots for illustration and are not performance measurements.

A. Dashboard and Navigation

After sign-in, a role-based dashboard (Fig. 3) summarises the selected organisation through eight indicators: number of clients, financial-year revenue and expenses, net profit with its margin, receivables, payables, closing cash, and pending approvals. A monthly revenue-versus-expense chart and an expense-distribution chart support quick review. The side navigation groups the modules as Clients; Finance (Financial Data, Financial Analysis, AI CFO, Forecasting, and Reports); Compliance (Overview, Compliance, and Approvals); and Governance (Agent Activity and Audit Logs). A persistent notice states that AI-generated analysis is not a substitute for a licensed CA or tax professional.

B. AI CFO Workspace and Forecasting

In the AI CFO workspace (Fig. 4), eight report tiles start the deterministic tools: Financial Summary, Profit and Loss, Cash Flow, Expense Analysis, Receivables, Payables, Budget versus Actual, and Monthly MIS. An ML Forecast panel plots the historical monthly series of a selected metric, revenue in the example, together with a dashed forecast for a selected horizon of three months. A chat panel lets the user ask the AI CFO to interpret a generated report. It states that the agent reads stored reports and that a report must be generated first to obtain exact figures, which follows the design rule that explanations rest on verified calculations.

C. Compliance Agent

The Compliance Agent page (Fig. 5) provides a statutory calendar, task tracking, and AI-assisted compliance status reports for a chosen organisation, month, and year. Counters show the total, upcoming, due-soon, overdue, and completed tasks, and an alert flags overdue filings. The Generate tasks action creates the statutory obligations for the month, covering GST, TDS, income tax, and ROC/MCA, and the Generate report action produces the status report. Each filing is listed with a status badge, its due date, and an action control; the example shows a TDS challan payment marked overdue. As in the CFO workspace, an assistant panel interprets the stored compliance tasks.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 10 , October - 2026

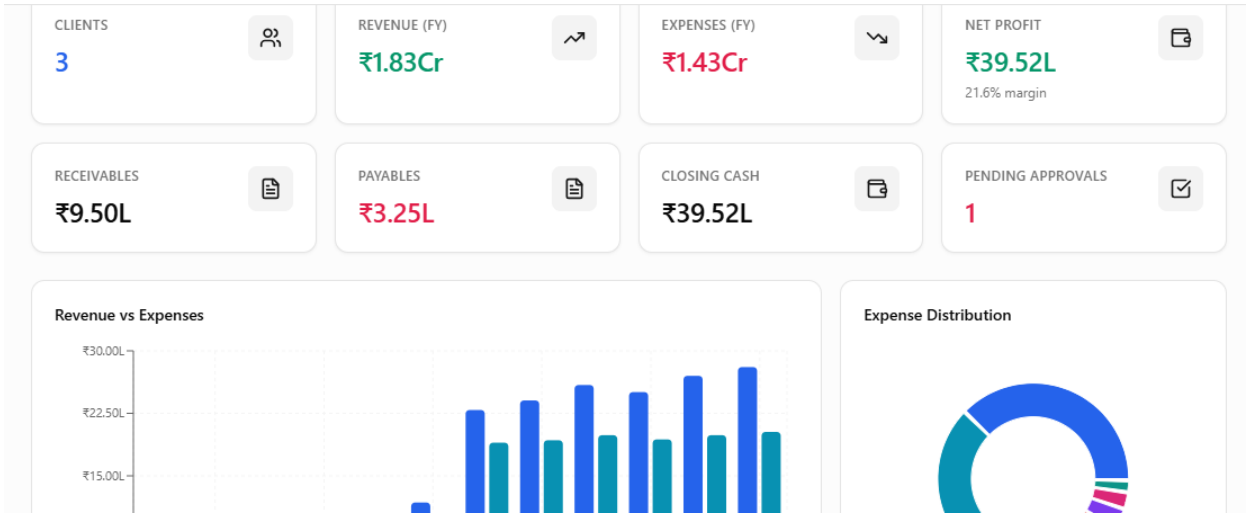


Fig. 3. FinVault AI dashboard showing eight financial indicators, the monthly revenue-versus-expense chart, and the expense distribution for an example organisation.

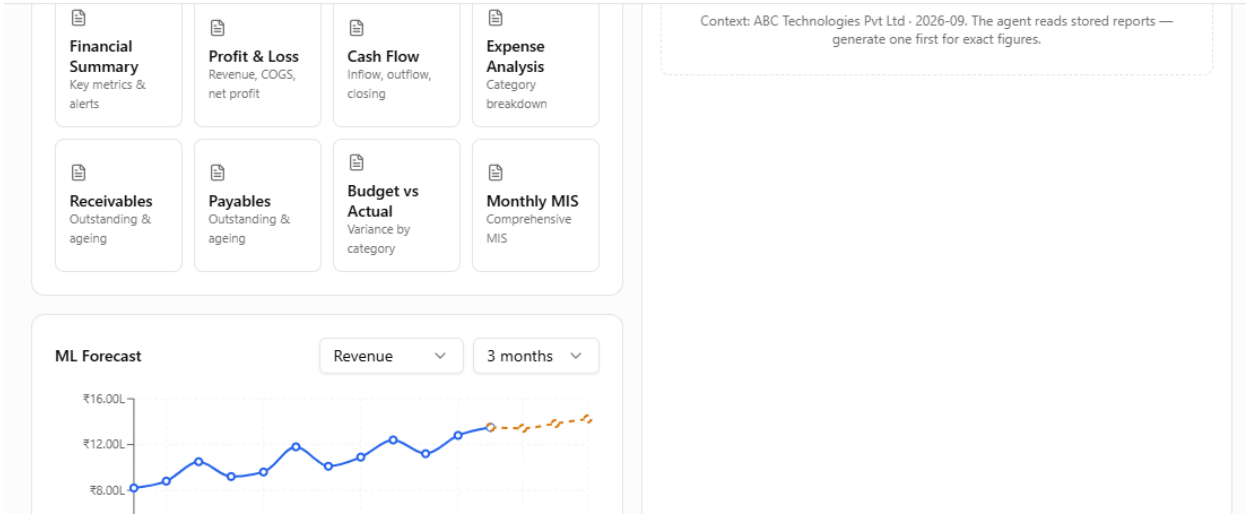


Fig. 4. AI CFO workspace with report tiles, the machine-learning forecast panel, and the assistant panel.

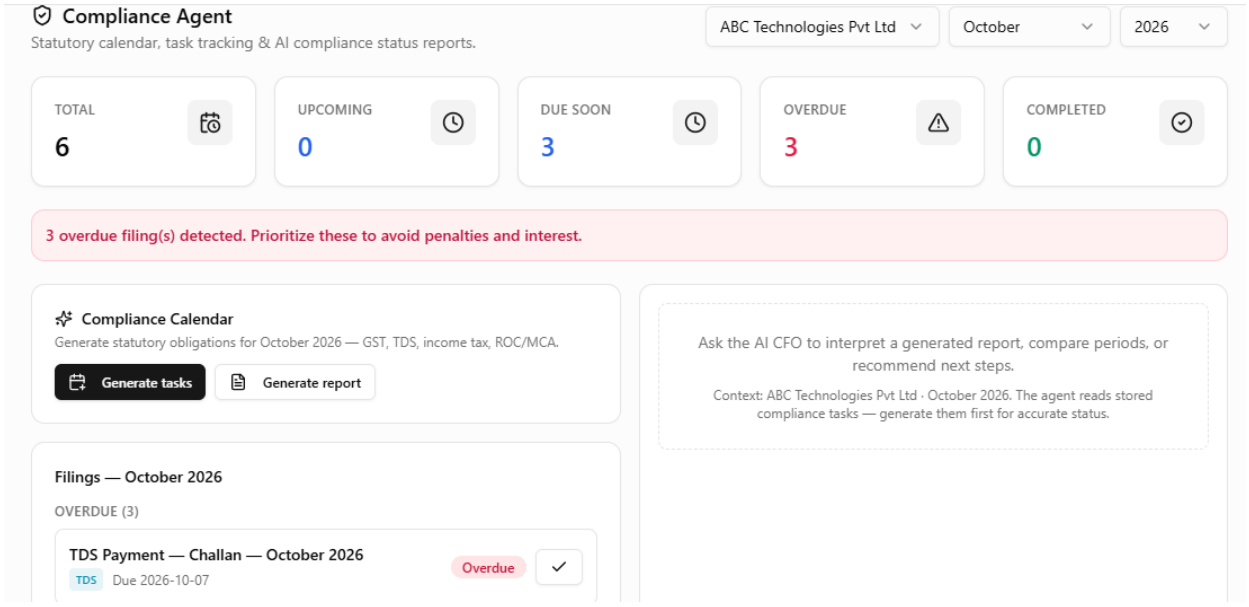


Fig. 5. Compliance Agent page with task counters, an overdue-filing alert, compliance calendar actions, and the filing list.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 10 , October - 2026

VIII. REQUIREMENTS AND SECURITY

A. Functional Requirements

The system shall:

- allow authorised users to manage clients and to upload and review financial data, transactions, and documents;
- calculate P&L, cash flow, receivables, payables, and budget variance with validated backend functions, and generate summaries and MIS reports;
- forecast revenue, expenses, and cash position when suitable historical data is available;
- track compliance tasks and deadlines, check required documents, identify missing items, and prepare client-wise reports;
- accept documents through WhatsApp, extract text and metadata, answer document questions by retrieval, and send expiry reminders where dates can be reliably identified; and
- hold sensitive outputs for human review and approval.

B. Non-Functional Requirements

- **Security:** authentication, authorisation, encrypted storage, secure data transfer, and tenant isolation.
- **Reliability:** dependable access to records, with backup and recovery.
- **Scalability:** support for growing numbers of users, clients, and documents, with long tasks run in the background.
- **Responsiveness:** low latency for common operations.
- **Transparency:** AI-generated explanations distinguishable from verified calculations, and important actions recorded in audit logs.
- **Maintainability:** modular services that can be updated independently, with errors and failed jobs monitored.

C. Security and Governance

Because the platform handles financial and personal data, security is designed into every layer. Users authenticate with JWT/OAuth 2.0 tokens, and role-based access control restricts each role to its permitted functions. Tenant isolation keeps the records of each organisation and client invisible to others, and documents and reports are held in access-controlled, encrypted storage. Audit logs record important actions, tool calls, agent runs, and approvals so that outputs can be traced.

Governance rules also limit the AI components. The LLM interprets requests and explains results but does not compute figures; tools run only under controlled permissions; AI-generated text is labelled apart from verified numbers; and sensitive outputs wait for CA approval. The system is limited to preparation and organisation. It does not make automatic government filings or tax payments, and it does not take fully autonomous financial decisions.

IX. METHODOLOGY AND EVALUATION PLAN

A. Development Methodology

The platform is developed iteratively in modular stages. First, user roles, workflows, data sources, and system boundaries are identified, keeping business workflows separate from personal DocVault workflows while identifying shared services. Second, the data model is designed, with tables for users, organisations, clients, transactions, documents, compliance tasks, approvals, reports, and audit events; documents are kept in object storage while their embeddings are indexed for retrieval. Third, document processing (validation, OCR where needed, metadata identification, and indexing) and the deterministic financial tools are implemented. Fourth, the agents and the approval workflow are built in LangGraph [8]. Finally, functional correctness, access control, data isolation, and response quality are tested.

B. Evaluation Plan

The deployed platform will be evaluated as summarised in Table III. Each candidate forecasting model will be compared with a simple baseline on held-out historical data using MAE, RMSE, and MAPE:

$$\text{MAE} = (1/n)\,\Sigma\,|y_i - \hat{y}_i| \quad (1)$$

$$\text{RMSE} = \sqrt{[(1/n)\,\Sigma\,(y_i - \hat{y}_i)^2]} \quad (2)$$

$$\text{MAPE} = (100\%/n)\,\Sigma\,|(y_i - \hat{y}_i)/y_i| \quad (3)$$

where the sums run over $i = 1, \ldots, n$; y_i is the observed value, $\hat{y}_i$ the predicted value, and n the number of observations.

TABLE III. PLANNED EVALUATION

Component	Planned check
Financial tools	Unit tests against hand-computed P&L, cash-flow, and variance values
Forecasting	MAE, RMSE, and MAPE against a baseline on held-out data
OCR	Field-level extraction accuracy on sample documents
RAG answers	Retrieval hit rate and agreement of answers with the source passage
Compliance engine	Correct detection of overdue tasks and missing documents
Security	Tests with invalid tokens, cross-tenant access, and unapproved reports
Background jobs	Completion time for OCR, indexing, and report generation

C. Expected Outcome

The platform is expected to reduce repetitive administrative effort and to improve access to organised financial, compliance, and document information. This expectation will be tested through the checks in Table III and through feedback from CAs and accountants, rather than assumed.

X. CONCLUSION AND FUTURE WORK

This paper presented the design and implementation of an AI-powered Virtual CFO and Compliance Automation Platform with an Intelligent Document Vault. By dividing the work among an AI CFO Agent, a Compliance Agent, and a DocVault Agent on shared infrastructure, the platform addresses repetitive financial processing, fragmented compliance tracking, and unsearchable personal documents. Its central design choice is to keep LLMs to interpretation and explanation, while deterministic tools perform calculations,

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 10 , October - 2026

OCR and RAG ground answers in stored records, and CA approval, access control, tenant isolation, and audit logs keep humans in charge of sensitive outputs. The web portal has been deployed as FinVault AI.

The work has limitations. No quantitative evaluation has been reported yet. Forecast quality depends on the amount and quality of historical data, and forecasts remain estimates. OCR accuracy varies with image quality, layout, language, and document condition, so extracted values need validation. LLM outputs can be wrong and must be checked against verified calculations and sources. The scope is an academic minimum viable product that excludes ERP replacement, automatic government filings, tax payments, and fully autonomous decisions.

Future work will carry out the evaluation in Table III on representative financial and document data, gather feedback from CAs and accountants, and extend the compliance rule set and the forecasting models.

ACKNOWLEDGMENT

The authors thank the Head of the Department, Dr. Vijaya Kumar K, and the Principal, Dr. Yuvaraju B. N., of AMC Engineering College, Bengaluru, for their support.

REFERENCES

[1] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, “Attention is all you need,” in *Advances in Neural Information Processing Systems*, vol. 30, 2017. [Online]. Available: https://arxiv.org/abs/1706.03762

[2] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, et al., “Language models are few-shot learners,” in *Advances in Neural Information Processing Systems*, vol. 33, 2020. [Online]. Available: https://arxiv.org/abs/2005.14165

[3] P. Lewis, E. Perez, A. Piktus, F. Petroni, V. Karpukhin, N. Goyal, H. Küttler, M. Lewis, W. Yih, T. Rocktäschel, S. Riedel, and D. Kiela, “Retrieval-augmented generation for knowledge-intensive NLP tasks,” in *Advances in Neural Information Processing Systems*, vol. 33, 2020. [Online]. Available: https://arxiv.org/abs/2005.11401

[4] R. Smith, “An overview of the Tesseract OCR engine,” in *Proc. 9th Int. Conf. Document Analysis and Recognition (ICDAR)*, 2007, doi: 10.1109/ICDAR.2007.4376991.

[5] S. Yao, J. Zhao, D. Yu, N. Du, I. Shafran, K. Narasimhan, and Y. Cao, “ReAct: Synergizing reasoning and acting in language models,” in *Proc. Int. Conf. Learning Representations (ICLR)*, 2023. [Online]. Available: https://arxiv.org/abs/2210.03629

[6] L. Breiman, “Random forests,” *Machine Learning*, vol. 45, no. 1, pp. 5–32, 2001, doi: 10.1023/A:1010933404324.

[7] T. Chen and C. Guestrin, “XGBoost: A scalable tree boosting system,” in *Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining*, 2016, pp. 785–794, doi: 10.1145/2939672.2939785.

[8] LangChain, “LangGraph documentation.” [Online]. Available: https://langchain-ai.github.io/langgraph/ (accessed Oct. 2026).

[9] S. Ramírez, “FastAPI documentation.” [Online]. Available: https://fastapi.tiangolo.com/ (accessed Oct. 2026).

[10] “pgvector: Open-source vector similarity search for Postgres.” [Online]. Available: https://github.com/pgvector/pgvector (accessed Oct. 2026).

(This work is licensed under a Creative Commons Attribution 4.0 International License.)