v. 5, n. 7, 2025

Journal of
Engineering
Research

Acceptance date: 24/09/2025

BLUEPRINT FOR INTELLIGENT MODERNIZATION OF LEGACY SYSTEMS: AI APPLIED TO THE FINANCIAL MARKET AND INCIDENT MANAGEMENT

Wagner Sabino dos Santos
Holds a Bachelor's degree in Educational Technology from Instituto Superior de Educação Ateneu – Vitória (2017), a Bachelor's degree in Information Systems from FAESA - Faculdades Integradas do Espírito Santo (2005), and a Postgraduate degree in Educational Technology from Instituto Superior de Educação Ateneu – Vitória. He is currently an SAP (Analytics Specialist Consultant) at Infraestruturas de Portugal

Abstract: This article presents a comparative case study between two legacy software modernization projects, applied in the financial market and IT service management sectors. Based on these cases, a robust and replicable methodological blueprint is proposed, aimed at the application of Artificial Intelligence (AI) and Machine Learning (ML) in consolidated platforms, with a focus on scalability, reuse, and architectural redesign. **The** Vera Cruz Connect Project demonstrated the feasibility of integrating Bloomberg and Refinitiv data with digital copilots, optimizing due diligence and risk analysis processes. The Mindworks Project, on the other hand, showed how intelligent automation in GLPI/ITSM environments can transform technical support into a core of operational intelligence. The results reveal significant reductions in operating costs, response time, and false positives, as well as a return on investment (ROI) of less than six months in both cases. The study consolidates itself as a practical and strategic reference by aligning technological innovation, governance, and cross-cutting applicability, positioning external AI as a central vector in the digital transformation of legacy systems.
Keywords: Blueprint, Redesign, Scalability, Reuse, Artificial Intelligence, Machine Learning, Digital Transformation, Technological Governance.

INTRODUCTION

The growing dependence on legacy systems in strategic sectors, such as the financial market and information technology service management (ITSM), highlights a recurring dilemma: how to modernize consolidated platforms without compromising their operational and regulatory reliability. Academic literature points out that the modernization of legacy software is often a costly and high-risk process, but inevitable to sustain scalability and competitiveness in rapidly changing digital environments (SOMMERVILLE, 2019; BASS; CLEMENTS; KAZMAN, 2021).

In this context, the application of Artificial Intelligence (AI) and Machine Learning (ML) has proven to be a viable strategy for "unlocking" value from existing data, creating intelligent layers that enhance efficiency and decision-making without the need to rewrite entire systems (RUSSELL; NORVIG, 2021). Recent literature in software engineering and data science confirms that the reuse of data and services through APIs (Application Programming Interface) and external integrations is a robust way to align innovation with governance and compliance (MARTIN; KUHN, 2020).

This article aims to propose and analyze a methodological blueprint for scalability and reuse in legacy systems, built from a comparison between two practical implementation case studies, located in distinct but complementary domains:

- Project 1: Vera Cruz Connect — implementation of a unified market intelligence platform for the Vera Cruz Group, with integration of Bloomberg and Refinitiv data, intelligent copilots, and automation of *due diligence* in international operations (Hong Kong, Dubai, and the US).
- Project 2: Mindworks ITSM with GLPI (Free IT Park Management) — intelligent automation of IT incidents in a Brazilian corporate environment, integrating AI and ML into the GLPI platform for ticket classification, prioritization, and forecasting, as well as interactive dashboards and proactive alerts.

Both cases illustrate how the "external AI" strategy — understood as the insertion of a layer of artificial intelligence via APIs, without changes to the core of the legacy software — enables the digital transformation of consolidated systems. This approach ensures scalability, reuse, and compliance, which are essential elements for organizations that simultaneously need innovation and operational stability (GARTNER, 2023).

Thus, this work is justified by its dual practical and academic contribution: on the one hand, it synthesizes real experiences of modernization in different domains (financial market and ITSM), demonstrating the cross-cutting applicability of the proposed model; on the other hand, it formalizes a replicable blueprint that can guide organizations and researchers in adopting similar strategies.

The structure of the article is organized as follows: after this introduction, the Theoretical Framework presents fundamental concepts about legacy software, scalability, reuse, and AI applications. Next, the Methodology describes the comparative analysis criteria. The Results and Discussion section presents the two projects individually, followed by a cross-analysis. Finally, the Conclusion summarizes the main findings, limitations, and directions for future research.

THEORETICAL FRAMEWORK

1. Legacy Software and Modernization Challenges: The term legacy software refers to systems that, although still indispensable to the functioning of organizations, were developed in outdated architectures, languages, or models, presenting limitations in terms of scalability, integration, and maintenance (SEACORD et al., 2016). According to Sommerville (2019), the central dilemma of these systems lies in the high cost of replacement versus the risk of functional and technological obsolescence.

Studies indicate that about 70% of the IT budget in large corporations is directed toward maintaining legacy systems, leaving little room for innovation (IDC, 2021). This scenario highlights the need for alternative approaches that allow for the incorporation of technological innovation without requiring the complete rewriting or replacement of software. In this context, modernization strategies based on additional layers of intelligence and integration have proven to be highly attractive.

2. Scalability and Reuse in Software Architectures: Scalability is defined as the ability of a system to maintain or improve its performance as it grows in data volume, users, or transactions (BASS; CLEMENTS; KAZMAN, 2021). Reuse, in turn, consists of adopting existing components, models, or services in different contexts or applications, reducing development effort and promoting consistency (PRESSMAN; MAXIM, 2021).

In this sense, modern architectures based on APIs and microservices have established themselves as effective alternatives for increasing scalability and reuse, allowing legacy systems to be extended without modifications to their core. This practical enables greater flexibility, facilitates integration with new technologies, and reduces dependence on rigid monolithic solutions (FOWLER, 2002).

3. Artificial Intelligence and Machine Learning in Consolidated Platforms: Artificial Intelligence (AI), defined by Russell and Norvig (2021) as the study of agents capable of perceiving and acting in environments in a rational manner, has found growing applications in legacy systems through external integration. The use of Machine Learning (ML), in particular, allows data already stored in these systems to be reinterpreted and enriched, generating new insights, predictions, and automations that were not previously possible.

According to Gartner (2023), the external AI strategy—building intelligent layers on top of existing systems—represents a pragmatic, scalable, and low-risk approach to digital modernization. In financial environments, for example, integrations with licensed APIs such as Bloomberg B-PIPE and Refinitiv RDP enable advanced risk analysis and due diligence automation (BLOOMBERG, 2023; LSEG, 2023). In the context of ITSM, platforms such as GLPI, when enriched by AI and ML, reach new levels of efficiency, with automated screening, call forecasting, and interactive dashboards, transforming technical support into a true operational intelligence center (SANTOS; MANCHESTER, 2025).

4. Methodological Blueprint as a Replicable Guide: The notion of blueprint in the field of software engineering refers to a methodological representation that guides the replication of practices in different scenarios (SCHMIDT; CLEMENS, 2019). By consolidating learnings from multiple cases into a structured model, the blueprint acts as a replicable guide, maximizing academic and practical impact.

In the present study, the blueprint emerges from the comparative analysis of two projects located in different domains (financial market and ITSM). This systematization seeks to coherently articulate the concepts of scalability, reuse, and AI applied to legacy systems, transforming practical experiences into a methodological model with the potential for generalization to other organizations and critical sectors.

METHODOLOGY

The research adopts a comparative case study as a methodological strategy, suitable for investigating complex phenomena in their real context and with multiple sources of evidence (YIN, 2015). The objective is to understand, based on two practical experiences in different environments, how the application of Artificial Intelligence (AI) and Machine Learning (ML) on legacy systems produces scalability and reuse of solutions, generating measurable transformations in critical operations.+

Study design.

Type: multiple case study, with replication logic (literal and theoretical) between the two cases.

Units of analysis: (i) analytical integration initiative with AI/ML in a financial context; (ii) automation and support initiative (ITSM) with AI/ML.

Guiding questions: (a) Which architectural and organizational decisions enable gains in scale/reuse? (b) Which metrics demonstrate such gains? (c) Which governance controls support the results?

Sources of evidence and collection.

- Technical documentation and project reports: reference architectures, executive plans, decision minutes, and impact reports.
- Operational metrics and logs: latency (average and p95), model accuracy/precision, TMR, automation rate, and cost savings; broker/API logs and dashboard time series.
- Semi-structured interviews: with analysts, managers, and technicians; recording, transcription, and double thematic coding. Data collection followed a standardized protocol (checklist by source, version, author, date) and triangulation between evidence to reinforce reliability (YIN, 2015).

Analytical procedures.

1. Data organization and preparation (field normalization, version control, traceability between artifacts).
2. Quantitative series analysis (moving windows and before/after comparison for KPIs).
3. Qualitative analysis with thematic coding and comparison of narratives between profiles.
4. Cross-case synthesis to identify design patterns, recurring decisions, and validity conditions.

Validity and reliability.

- Construct validity: multiple sources and documented chain of evidence.
- Internal validity: explanation of plausible causal links between decisions and results; analysis of threats (bias, data maturity).
- External validity: logic of replication between cases and rich description of the context for transferability.
- Reliability: use of study protocol and versioned repository of data/artifacts (audit trail), as recommended by Yin (2015).

Ethical and compliance aspects.

The data followed principles of minimization and pseudonymization; interviews were conducted with free and informed consent; access to repositories complied with segregation of duties (SoD) and internal retention policies (according to LGPD).

Success criteria.

Robust evidence of impact was considered when: (i) there was consistent improvement in technical/business KPIs (e.g., p95 latency, TMR, false positives, savings); (ii) reuse was demonstrated by versioned artifacts applicable in more than one flow; and (iii) the results showed sustainability under governance and operational audit.

CHARACTERIZATION OF CASES

Case 1 – Vera Cruz Connect (Financial Market): The Vera Cruz Group, a multinational conglomerate active in private equity, asset management, and commodities, faced challenges related to real-time market data ingestion, due diligence execution, and compliance with regulatory audits in multiple jurisdictions.

The project consisted of creating a unified market intelligence platform, integrating data from Bloomberg (B-PIPE/BLPAPI) and Refinitiv (RDP/EDAPI), applying AI for risk analysis, and digital co-pilots to support analysts and investors. The architecture was designed to operate in multiple regions (Hong Kong, Dubai, and the US), prioritizing compliance, scalability, and reuse across the group's different units.

Case 2 – Mindworks ITSM with GLPI (Incident Management): Mindworks, a Brazilian IT solutions company, had bottlenecks in incident management on its GLPI platform, characterized by manual screening, low predictability, and a lack of strategic insights for technical support.

The project sought to transform support into an operational intelligence center through automation of triage, call forecasting, and real-time dashboard generation. The architecture implemented used GLPI REST API connectors and ML algorithms (Scikit-learn, Prophet, XGBoost), resulting in reduced operating costs and significant improvement in SLA indicators.

COMPARISON CRITERIA AND METRICS

The comparative analysis was structured based on quantitative and qualitative criteria widely recognized in the literature on software engineering and IT management (PRESSMAN; MAXIM, 2021; BASS; CLEMENTS; KAZMAN, 2021):

- Scalability
  - o Volume of events/second processed.
  - o Growth capacity for multiple users/assets.
  - o Latency in critical data ingestion.
- Reusability
  - o Degree of solution modularity.
  - o Ease of replication in new environments.
  - o Adaptation to different regulations or internal processes.

Journal of Engineering Research ISSN 2764-1317 DOI https://doi.org/10.22533/at.ed.3175725170716

- Operational Performance

o Average response time (ART).

o Task automation rate (e.g., ticket screening, due diligence).

o Reduction of false positives (in compliance or incident classification).

- Business Impact

o Return on investment (ROI).

o Reduction in operating costs.

o Efficiency gains (audits, onboarding, technical support).

- Quality and Governance

o Regulatory compliance (financial, GDPR, SLA).

o Explainability of AI models (explainable AI).

o End-user satisfaction (analysts, technicians, customers).

Data Collection Procedures

The data used in this study was obtained from three complementary sources:

- Technical Documentation and Project Reports

o Reference architectures, executive plans, and impact reports were analyzed with version, author, and date metadata, ensuring traceability between decisions and evidence provided by technical teams.

- Performance Metrics and Operational Logs

o Latency (average and p95), model accuracy/precision, TMR, automation rate, and cost savings were measured.

o Logs were collected from Kafka (throughput, errors per topic), GLPI API (ticket cycle), and Power BI/Python Dash time series (CSV/Parquet export), with file version control.

- Qualitative Context Analysis

o Semi-structured interviews were conducted with free and informed consent, with double coding of transcripts and resolution of discrepancies by consensus, with stakeholders (financial analysts, risk managers, IT technicians).

o Observation of governance and compliance practices in each domain. For reliability, checks were applied for completeness, temporal consistency, and congruence between systems; for ethics and compliance, pseudonymization, data minimization, and segregation of duties (SoD) were adopted, in addition to retention in accordance with internal policy and LGPD.

STUDY DESIGN

The methodological design can be summarized in four sequential steps:

- Project characterization — detailed description of the objectives, solutions, and results achieved in each case.
- Extraction of metrics — identification and standardization of key indicators comparable between the two projects.
- Cross-analysis — comparison of results against defined criteria, highlighting convergences and divergences.
- Formalization of the blueprint — synthesis of replicable practices and lessons learned, composing the methodological guide for legacy software modernization.

COMPARATIVE ANALYSIS AND DISCUSSION

Cross-analysis of the two projects allows for the identification of relevant patterns and contrasts, which consolidate the proposed methodological blueprint.

Scalability:

- Scalability was evident in each case: at the Vera Cruz Group, through the ability to process a high volume of events in real time and with low latency, which is essential for the financial market. At Mindworks, through the ability to manage many assets and tickets, which

Criterion	Vera Cruz Project (Financial)	Mindworks Project (ITSM)	Comparative Analysis
Scalability	Processing of 100k+ events/s, latency <250ms	Scalability for +10,000 monitored assets	Both demonstrated robust scalability, albeit in different domains (real-time financial data vs. corporate ticket management).
Reuse	Modular architecture replicated in Hong Kong, Dubai, and the US	Rapid replication via GLPI API in multiple instances	Vera Cruz demonstrated multinational reuse; Mindworks demonstrated technical replicability across multiple clients/ environments.
Operational Performance	-45% analysis time, -35% AML false positives	-42% TMR, +72% SLA compliance	Both reduced critical times: market analysis and incident response.
Business Impact	ROI < 6 months, savings >USD 500k/ year	ROI < 5 months, savings R$ 190k/year	Both achieved accelerated ROI, validating the economic viability of the approach.
Quality/Governance	Explainable AI, auditable compliance	Monitored SLA, ticket traceability	In both cases, AI was applied with a focus on transparency and trust, a critical factor for organizational adoption.

is crucial for a growing ITSM environment.

- The use of external APIs, Kafka, and Data Lakehouse enables this scalability in such distinct domains. Emphasizing the ability to handle growing volumes of data and requests, it has become a fundamental pillar of the modernization model.

Reuse:

- At the Vera Cruz Group, the modular architecture allowed replication in different regions (Hong Kong, Dubai, USA), demonstrating the ability to adapt to different regulatory and operational contexts. At Mindworks, reuse occurred in the ability to apply the same base solution (via GLPI API) in multiple instances or clients, demonstrating the technical replicability of the model.
- Modularity, the use of APIs for integration, and the abstraction of the AI layer allowed the solution to be quickly adapted and deployed in new scenarios, reducing the time and cost of development for each new application.

Operational Performance:

- At the Vera Cruz Group, the reduction in analysis time and AML/KYC false positives directly impacted the agility of investment decisions and compliance. At Mindworks, the reduction in TMR and increased SLA compliance improved support efficiency and end-user satisfaction.
- The operational gains are not merely incremental, but transformational, driven by the predictive and automated capabilities of AI/ML. The blueprint provides a framework for identifying and optimizing operational bottlenecks in different types of systems.

Business Impact:

- The rapid ROI (less than 6 months) in such diverse contexts validates the blueprint's value proposition. The savings in man-hours (reallocation of FTEs at Vera Cruz) and reduction in operating costs (Mindworks) demonstrate the tangible value of AI-driven modernization.
- It is not just a technical solution, but a business strategy, allowing organizations not only to modernize their systems, but also to generate significant financial value in a short period of time, making the adoption of AI a strategic and economically viable decision.

Quality/Governance:

- At the Vera Cruz Group, explainable AI and auditable compliance are crucial for a highly regulated sector such as finance. At Mindworks, SLA monitoring and ticket traceability increase transparency and confidence in service management.
- The transparency, trust, and auditability of AI is an integral component of the blueprint. This makes it vital for adoption in critical corporate environments, where AI cannot be a "black box," but rather a reliable and accountable tool.

DISCUSSION

Both cases confirm that value was not created from scratch, but unlocked from data already existing on legacy platforms. Both Bloomberg/Refinitiv and GLPI contained rich but underutilized information. The introduction of the external AI layer made it possible to extract new meanings, automate processes, and increase the reliability of operations.

Another common point is the strategic impact. While in the case of the Vera Cruz Group, AI enabled faster and safer global financial decisions, in the Mindworks case, it transformed IT into an engine of operational efficiency. In both cases, AI acted as a catalyst for the transformation of critical operations.

The comparative analysis also shows that the versatility of the blueprint does not depend on the application domain: the same logic of modular architecture, integration via APIs, and ML models can be adapted to regulated (financial) and corporate (ITSM) sectors. This reinforces the method's replicability as a structured reference for the modernization of legacy systems.

CONCLUSION

The comparative analysis of the Vera Cruz Connect and Mindworks ITSM projects with GLPI indicates that the modernization of legacy systems through external layers of AI is an effective strategy for scalability, reuse, and transformation of critical operations. In both cases, the hypothesis that value does not need to be created from scratch, but rather unlocked from data already present in consolidated platforms, expanding its usefulness and impact, was confirmed.

At Grupo Vera Cruz, integration with highly critical financial data (Bloomberg, Refinitiv) and the creation of analytical co-pilots have enabled efficiency gains in market analysis, due diligence automation, and increased regulatory reliability. In the case of Mindworks, intelligent incident automation in GLPI transformed technical support into a strategic intelligence center, with a direct impact on cost reduction, increased SLA compliance, and strengthened IT governance.

From a practical standpoint, this study presented evidence that modular architectures, based on APIs and external AI layers, are capable of delivering measurable results in short timeframes, with ROI in less than six months. From an academic standpoint, systematizing these learnings into a methodological blueprint expands the literature on legacy software modernization, offering a replicable guide that can be applied across different industries.

Study Limitations: Despite the natural limitations of the scope, the results obtained demonstrate strong potential for scalability and applicability of the proposed model.

The analysis focused on two case studies, which limits the statistical generalization of the findings. However, the technical depth and diversity of the contexts analyzed provide a solid basis for replication in similar environments.

The presence of distinct sectors—financial market and ITSM—enriches the comparability between domains with different operational and technological requirements. Still, caution is recommended when extrapolating the results to sectors with substantially different dynamics.

Some performance and impact metrics were provided by the companies themselves, which may introduce measurement or reporting biases. However, the data were treated with technical criteria and triangulated with operational evidence, reinforcing the reliability of the conclusions.

Future Perspectives: The results obtained open promising paths for the consolidation and expansion of the proposed model. As strategic developments, the following are suggested:

- Sectoral expansion of the blueprint to critical areas such as health, energy, and logistics, aiming to validate its cross-cutting applicability in environments with high regulatory, operational, and technological complexity. The proposed architecture demonstrates sufficient flexibility to adapt to different domains, reinforcing its potential as a methodological reference.
- Longitudinal studies are recommended to assess the sustainability of scalability and reuse gains over time, especially in scenarios of continuous evolution of requirements, technologies, and data volumes.
- Comparative analyses with alternative approaches, such as complete system rewriting or direct migration to SaaS platforms, may reveal relevant trade-offs between cost, implementation time, operational risks, and strategic impact. This comparison is essential to guide decisions in contexts with specific constraints.
- In-depth investigations into AI ethics, governance, and explainability, especially in regulated industries, where the traceability of automated decisions is not only desirable but mandatory. The integration of transparency and compliance mechanisms can strengthen institutional and social acceptance of AI-based solutions.

This work consolidates the notion that the modernization of legacy software through artificial intelligence is technically feasible, strategically transformative, and methodologically replicable. By formalizing a scalable and reusable blueprint, it contributes to the advancement of software engineering and data science, providing a solid and applicable reference for researchers, system architects, and innovation leaders committed to the convergence of technology, regulatory compliance, and real value generation.

REFERENCES

BASS, L.; CLEMENTS, P.; KAZMAN, R. Software Architecture in Practice. 4. ed. Boston: Addison-Wesley, 2021.

BLOOMBERG. Bloomberg API and B-PIPE Documentation. 2023. Disponível em: https://data.bloomberglp.com. Acesso em: 3 set. 2025.

FOWLER, M. Patterns of Enterprise Application Architecture. Boston: Addison-Wesley, 2002.

GARTNER. Artificial Intelligence and Emerging Technologies: External AI Strategies. Stamford, 2023.

IDC. Worldwide IT Spending Guide: Legacy Systems and Modernization. Framingham: IDC, 2021.

LSEG – London Stock Exchange Group. Refinitiv Data Platform APIs: Developer Documentation. 2023. Disponível em: https://developers.lseg.com. Acesso em: 3 set. 2025.

MARTIN, R. C.; KUHN, S. Clean Architecture: A Craftsman's Guide to Software Structure and Design. Boston: Prentice Hall, 2017.

PRESSMAN, R. S.; MAXIM, B. R. Engenharia de Software: Uma Abordagem Profissional. 9. ed. Porto Alegre: AMGH, 2021.

RUSSELL, S.; NORVIG, P. Artificial Intelligence: A Modern Approach. 4. ed. Harlow: Pearson, 2021.

SCHMIDT, D. C.; CLEMENS, R. Software Blueprinting and Reuse Patterns. New York: Springer, 2019.

SEACORD, R.; et al. Modernizing Legacy Systems: Software Architecture Perspectives. Upper Saddle River: Addison-Wesley, 2016.

SOMMERVILLE, I. Engenharia de Software. 10. ed. São Paulo: Pearson Education, 2019.

YIN, R. K. Case Study Research and Applications: Design and Methods. 6. ed. Thousand Oaks: Sage, 2015.