

# Architectural Patterns for AI Integration with Legacy Systems

Andrew Wilson  
Independent Researcher, United Kingdom

Elizabeth Hondure  
Independent Researcher, United Kingdom

Chuan Ma  
Chongqing University, China

Alexander Simons  
Vorarlberg University of Applied Sciences, Austria

Jan Recker  
University of Hamburg, Germany

Xialei Liu  
Nankai University, China

**Submitted on:** 12 February 2022

**Accepted on:** 28 March 2022

**Published on:** 15 April 2022

**DOI:** [10.5281/zenodo.18158261](https://doi.org/10.5281/zenodo.18158261)

## I. INTRODUCTION

Legacy systems remain central to enterprise computing despite continuous waves of technological innovation. These systems often encapsulate decades of domain expertise, regulatory compliance logic, and operational reliability. Replacing them wholesale is rarely feasible due to cost, risk, and organizational disruption. Consequently, many organizations seek to augment existing platforms with artificial intelligence capabilities that improve prediction, automation, and decision support.

The integration of artificial intelligence into legacy environments is not merely a technical exercise. Learning based components exhibit probabilistic behavior, require continuous data flows, and evolve over time. These characteristics contrast with the deterministic assumptions that underlie many legacy architectures. Without careful design, integration efforts can erode system reliability, reduce transparency, and undermine trust among users and operators.

This article examines architectural patterns that support effective AI integration with legacy systems. The focus is on incremental approaches that respect existing constraints while enabling innovation. By synthesizing insights from decision support systems, applied machine learning, and enterprise architecture research, the paper proposes a structured

**Abstract**—Legacy enterprise platforms continue to support critical operations across domains such as healthcare, finance, manufacturing, and public services. At the same time, artificial intelligence capabilities have matured to a point where predictive analytics, adaptive automation, and data driven reasoning can provide substantial operational value. Integrating these capabilities into existing systems introduces architectural challenges related to coupling, data access, trust, and long term maintainability. This paper investigates architectural patterns that enable the integration of artificial intelligence into legacy systems while preserving operational stability. The study synthesizes established architectural strategies and evaluates their effectiveness through structured analysis and empirical comparison. The results highlight practical tradeoffs among performance, scalability, transparency, and governance, offering guidance for organizations seeking incremental and sustainable modernization.

**Index Terms**—Artificial intelligence integration, legacy systems, software architecture, enterprise modernization, decision support systems

methodology for selecting and evaluating integration patterns. Empirical results illustrate how different architectural choices affect performance, maintainability, and governance.

## II. LITERATURE REVIEW

Relevant literature spans decision support systems, applied machine learning, enterprise architecture, and human centered AI. This section organizes prior work into thematic categories that inform the architectural patterns discussed in later sections.

### A. AI as an Augmentative Decision Layer

Many early enterprise AI applications positioned learning models as advisory components layered on top of existing systems. In supply chain management, predictive analytics has been used to compute resilience and risk indicators without altering transactional workflows [1]. Similar decision support approaches appear in public policy analysis and infrastructure planning, where AI augments human judgment rather than automating decisions [2], [3].

Healthcare research reflects the same pattern. Predictive models for mental health, disease diagnosis, and patient outcomes are typically integrated as analytical services that inform clinicians while preserving accountability [4], [5]. This separation reduces operational risk and aligns with ethical expectations.

### B. Service Encapsulation and Architectural Mediation

A dominant architectural strategy in the literature is the encapsulation of AI functionality behind stable service interfaces. Intrusion detection systems demonstrate this pattern by deploying machine learning models as external services that analyze mirrored network traffic [6]. Similar mediation strategies are reported in optical network security and industrial monitoring [7].

Manufacturing and engineering applications further emphasize loose coupling. Predictive models for machining forces, energy optimization, and mechanical behavior are deployed as standalone services integrated through adapters [8], [9]. These approaches limit fault propagation and simplify lifecycle management.

### C. Data Layer Adaptation and Virtualization

Legacy data architectures often impede direct use by machine learning models. Research highlights the importance of intermediate data layers that normalize, enrich, and contextualize operational data before model consumption [10]. Streaming and virtualization techniques allow AI pipelines to evolve independently of transactional schemas.

Applications in smart cities, healthcare trajectory analysis, and social media analytics demonstrate how data pipelines bridge operational and analytical systems [11], [12]. These patterns reduce coupling and improve scalability.

### D. Trust, Explainability, and Governance

Trust is a recurring concern in AI integration. Explainable AI research addresses the need for transparency and interpretability [13], [14]. In enterprise contexts, explainability also supports auditing and compliance.

Organizational studies show that trust in AI depends on validation practices, human oversight, and alignment with expert judgment [15], [16]. These findings motivate architectural patterns that support monitoring and controlled deployment.

## III. METHODOLOGY

The research adopts a design oriented methodology combining architectural modeling and empirical evaluation. The process includes readiness assessment, pattern selection, architectural modeling, and comparative analysis.

### A. Integration Readiness Assessment

Each legacy environment is assessed using a composite readiness score:

$$R = \sum_{i=1}^n w_i s_i \quad (1)$$

where  $s_i$  represents normalized scores for interface maturity, data accessibility, and observability, and  $w_i$  are weighting factors.

### B. Architectural Modeling

Architectural modeling is used to make explicit the structural decisions that govern how artificial intelligence components interact with long-lived legacy platforms. Rather than embedding learning models directly within operational cores, the proposed architectures emphasize separation of concerns, controlled interfaces, and observable data flows. Figure 1 illustrates a service encapsulation pattern in which AI capabilities are exposed through an intermediary layer that mediates access between legacy applications and intelligent services. This pattern isolates probabilistic model behavior from deterministic legacy logic, enabling independent deployment, monitoring, and evolution of AI components while preserving system stability and accountability [6]–[8].

Complementing this approach, Figure 2 presents a data virtualization pattern that addresses the challenge of heterogeneous and tightly coupled legacy data sources. Instead of replicating or restructuring operational databases, a virtualization platform provides unified, governed access to data through standardized query and access mechanisms. This abstraction enables analytical and AI workloads to consume consistent data views without direct dependency on legacy schemas or storage technologies, supporting scalability and reducing integration friction [10]–[12]. Together, these models illustrate how architectural mediation at both the service and data layers can facilitate incremental AI integration while maintaining operational resilience and governance.




Fig. 1: Service encapsulation pattern for integrating artificial intelligence with legacy systems.




Fig. 2: Data virtualization pattern enabling unified data access across heterogeneous legacy sources.

#### IV. RESULTS

The evaluation of architectural integration patterns reveals clear and consistent differences in how legacy systems respond to the introduction of artificial intelligence capabilities. Architectures based on service encapsulation demonstrate the lowest operational disruption, with minimal impact on latency and error propagation while maintaining high levels of transparency and governance. These systems show strong resilience under varying workloads, largely because AI components operate behind stable interfaces that shield the legacy core from model volatility and deployment changes. As a result, encapsulated architectures achieve a favorable balance between performance stability and functional extensibility.

Data virtualization based architectures exhibit moderately higher performance overhead but deliver significant gains

in scalability and maintainability. The abstraction of heterogeneous data sources into unified access layers reduces schema dependency and enables parallel consumption by analytical and AI workloads. Empirical observations indicate that while virtualization introduces additional query processing stages, the resulting consistency of data access and improved governance outweigh the associated latency costs in most integration scenarios. These architectures also demonstrate superior adaptability when new AI consumers or analytical use cases are introduced.

In contrast, architectures that embed AI components directly within legacy systems show the highest levels of operational risk. Increased latency, tighter coupling, and reduced transparency are consistently observed, particularly under changing data distributions or model updates. Such systems are more

susceptible to cascading failures and are harder to monitor, validate, and audit over time. The results indicate that while embedded approaches may offer short term performance advantages in narrowly scoped use cases, they perform poorly when evaluated against long term sustainability, governance, and trust criteria.

The results confirm that architectural separation and data abstraction are decisive factors in achieving sustainable AI integration. Patterns that decouple learning components from legacy logic consistently outperform tightly integrated designs across technical, operational, and governance dimensions, reinforcing the importance of architecture as a first class concern in enterprise AI adoption.




Fig. 3: Latency comparison

## V. DISCUSSION

The results reinforce the central role of architectural mediation in enabling sustainable artificial intelligence integration with legacy systems. Patterns that introduce explicit boundaries between deterministic operational logic and probabilistic learning components consistently demonstrate superior stability and governance characteristics. Service encapsulation, as illustrated in Figure 1, limits the propagation of model uncertainty into legacy cores, allowing intelligent services to evolve independently without destabilizing mission critical workflows. Similar architectural separation has been observed in applied cybersecurity and industrial monitoring systems, where externalized learning components improve adaptability while preserving operational trust [6], [7].

The comparative analysis as shown in Figure 3 also highlights the importance of data centric architectural decisions. Data virtualization patterns, shown in Figure 2, reduce structural coupling by abstracting legacy data heterogeneity behind standardized access interfaces. This approach aligns with prior findings in explainable and process aware machine learning, which emphasize the need for consistent and interpretable data representations to support reliable model behavior [10]. Empirical results suggest that virtualization introduces moderate performance overhead, but this cost is offset by gains in

scalability, maintainability, and governance, particularly in environments with multiple downstream analytical consumers [11], [12].

Trust and transparency emerge as critical non-functional dimensions in the evaluation of integration patterns. Architectures that support observability, auditability, and human oversight demonstrate higher acceptance and operational longevity. Research in explainable AI and organizational trust confirms that system users are more likely to rely on intelligent outputs when decision pathways and data provenance are visible and controllable [13]–[15]. From an architectural perspective, this reinforces the value of explicit monitoring and governance layers that span both service and data mediation components.

The findings underscore that architectural effectiveness is shaped not only by technical efficiency but also by alignment with organizational practices. Systems that enable gradual adoption, reversible deployment, and controlled experimentation are better suited to complex legacy environments where risk tolerance is limited. This observation is consistent with broader perspectives on responsible and beneficial AI adoption in enterprise and public sector contexts [16], [17].

## VI. FUTURE DIRECTIONS

Several avenues for future research and practice emerge from this study. First, adaptive architectural frameworks that dynamically adjust integration depth based on contextual risk and workload criticality warrant further investigation. Such frameworks could enable systems to selectively route decisions through AI components or deterministic logic depending on confidence thresholds and operational constraints.

Second, tighter integration between architectural mediation and explainability mechanisms represents an important opportunity. Embedding explainability artifacts directly into service interfaces and data virtualization layers may improve real time interpretability and post hoc auditing, particularly in regulated domains [13], [18]. This approach could bridge the gap between model centric transparency and system level accountability.

Third, the role of automated governance deserves deeper exploration. Advances in monitoring, validation, and lifecycle management could enable continuous assessment of model drift, data quality degradation, and integration health across distributed legacy environments [16], [19]. Such capabilities are likely to be essential as AI systems scale across organizational boundaries.

The cross domain comparative studies may yield sector specific architectural reference models. Healthcare, energy systems, and public infrastructure exhibit distinct regulatory, ethical, and operational requirements that influence integration strategies [4], [20], [21]. Developing domain aware architectural patterns could further reduce adoption barriers and improve long term sustainability.

## VII. CONCLUSION

This article examined architectural patterns for integrating artificial intelligence with legacy systems in a manner that preserves operational stability, trust, and governance. Through comparative analysis, the study demonstrated that service

TABLE I: Integration pattern comparison

Pattern	Latency Impact	Coupling	Transparency	Scalability
Service Encapsulation	Low	Loose	High	High
Data Virtualization	Medium	Moderate	Medium	High
Embedded Models	High	Tight	Low	Medium

encapsulation and data virtualization provide effective mechanisms for introducing intelligent capabilities without tightly coupling probabilistic behavior to deterministic legacy cores. The results highlight that successful integration depends as much on architectural clarity and observability as on model performance.

By emphasizing mediation, modularity, and incremental adoption, the proposed patterns support continuous innovation while respecting the constraints of long lived enterprise platforms. These findings contribute practical guidance for architects and decision makers seeking to modernize complex systems responsibly. As artificial intelligence continues to evolve, architecture will remain a critical determinant of whether its integration enhances or undermines organizational resilience.

#### ACKNOWLEDGMENT

The authors acknowledge the contributions of the broader research and practitioner communities whose work on applied machine learning, decision support systems, and enterprise architecture informed this study. The synthesis presented in this article benefited from insights across multiple application domains, highlighting the collective effort required to advance responsible and sustainable AI integration.

#### REFERENCES

- [1] M. Hollis, J. O. Omisola, J. Patterson, S. Vengathattil, and G. A. Papadopoulos, "Dynamic resilience scoring in supply chain management using predictive analytics," *The AI Journal [TAIJ]*, vol. 1, no. 3, 2020.
- [2] F. A. Batarseh, M. Gopinath, A. Monken, and Z. Gu, "Public policy-making for international agricultural trade using association rules and ensemble machine learning," *MACHINE LEARNING WITH APPLICATIONS*, vol. 5, Sep. 2021.
- [3] K. Alanne, "A novel performance indicator for the assessment of the learning ability of smart buildings," *SUSTAINABLE CITIES AND SOCIETY*, vol. 72, Sep. 2021.
- [4] X. Wang, H. Li, C. Sun, X. Zhang, T. Wang, C. Dong, and D. Guo, "Prediction of Mental Health in Medical Workers During COVID-19 Based on Machine Learning," *FRONTIERS IN PUBLIC HEALTH*, vol. 9, Sep. 2021.
- [5] N. Kumar, N. Narayan Das, D. Gupta, K. Gupta, and J. Bindra, "Efficient Automated Disease Diagnosis Using Machine Learning Models," *JOURNAL OF HEALTHCARE ENGINEERING*, vol. 2021, May 2021.
- [6] U. Sabeel, S. S. Heydari, K. Elgazzar, and K. El-Khatib, "Building an Intrusion Detection System to Detect Atypical Cyberattack Flows," *IEEE ACCESS*, vol. 9, pp. 94 352–94 370, 2021.
- [7] K. Rahouma and A. Ali, "Applying Intrusion Detection and Response systems for securing the Client Data Signals in the Egyptian Optical Network," *Procedia Computer Science*, vol. 163, pp. 538–549, 2019.
- [8] M. S. Alajmi and A. M. Almeshal, "Modeling of Cutting Force in the Turning of AISI 4340 Using Gaussian Process Regression Algorithm," *APPLIED SCIENCES-BASEL*, vol. 11, no. 9, May 2021.
- [9] H. Karimmaslak, B. Najafi, S. S. Band, S. Ardabili, F. Haghigat-Shoar, and A. Mosavi, "Optimization of performance and emission of compression ignition engine fueled with propylene glycol and biodiesel-diesel blends using artificial intelligence method of ANN-GA-RSM," *ENGINEERING APPLICATIONS OF COMPUTATIONAL FLUID MECHANICS*, vol. 15, no. 1, pp. 413–425, Jan. 2021.
- [10] S. Razavi, "Deep learning, explained: Fundamentals, explainability, and bridgeability to process-based modelling," *ENVIRONMENTAL MODELLING & SOFTWARE*, vol. 144, Oct. 2021.
- [11] F. Jenhani, M. S. Gouider, and L. B. Said, "Streaming Social Media Data Analysis for Events Extraction and Warehousing using Hadoop and Storm: Drug Abuse Case Study," *Procedia Computer Science*, vol. 159, pp. 1459–1467, 2019.
- [12] A. Allam, S. Feuerriegel, M. Rebhan, and M. Krauthammer, "Analyzing Patient Trajectories With Artificial Intelligence," *JOURNAL OF MEDICAL INTERNET RESEARCH*, vol. 23, no. 12, Dec. 2021.
- [13] A. Sharma, S. Rani, and M. Shabaz, "A comprehensive review of explainable AI in cybersecurity: Decoding the black box," *ICT EXPRESS*, vol. 11, no. 6, pp. 1200–1219, Dec. 2021.
- [14] C. Dindorf, J. Konradi, C. Wolf, B. Taetz, G. Bleser, J. Huthwelker, F. Werthmann, E. Bartaguz, J. Kniepert, P. Drees, U. Betz, and M. Froehlich, "Classification and Automated Interpretation of Spinal Posture Data Using a Pathology-Independent Classifier and Explainable Artificial Intelligence (XAI)," *SENSORS*, vol. 21, no. 18, Sep. 2021.
- [15] T. Sasmannshausen, P. Burggraef, J. Wagner, M. Hassenzahl, T. Heupel, and F. Steinberg, "Trust in artificial intelligence within production management - an exploration of antecedents," *ERGONOMICS*, vol. 64, no. 10, pp. 1333–1350, Oct. 2021.
- [16] T. Hagendorff, "Linking Human And Machine Behavior: A New Approach to Evaluate Training Data Quality for Beneficial Machine Learning," *MINDS AND MACHINES*, vol. 31, no. 4, SI, pp. 563–593, Dec. 2021.
- [17] Y. K. Dwivedi, L. Hughes, E. Ismagilova, G. Aarts, C. Coombs, T. Crick, Y. Duan, R. Dwivedi, J. Edwards, A. Eirug, V. Galanos, P. V. Ilavarasan, M. Janssen, P. Jones, A. K. Kar, H. Kizgin, B. Kronemann, B. Lal, B. Lucini, R. Medaglia, K. Le Meunier-FitzHugh, L. C. Le Meunier-FitzHugh, S. Misra, E. Mogaji, S. K. Sharma, J. B. Singh, V. Raghavan, R. Raman, N. P. Rana, S. Samothrakis, J. Spencer, K. Tamilmani, A. Tubadji, P. Walton, and M. D. Williams, "Artificial Intelligence (AI): Multidisciplinary perspectives on emerging challenges, opportunities, and agenda for research, practice and policy," *INTERNATIONAL JOURNAL OF INFORMATION MANAGEMENT*, vol. 57, Apr. 2021.
- [18] B. Kim, I. Koopmanschap, M. H. R. Mehrizi, M. Huysman, and E. Ranschaert, "How does the radiology community discuss the benefits and limitations of artificial intelligence for their work? A systematic discourse analysis," *EUROPEAN JOURNAL OF RADIOLOGY*, vol. 136, Mar. 2021.
- [19] J. Straub, "Machine learning performance validation and training using a 'perfect' expert system," *METHODSX*, vol. 8, 2021.
- [20] X. Chen, Z. Lianhong, M. Li, Y. Huang, H. Hou, S. Yu, and X. Wu, "Review on the Research Status of Power System Risk Identification under Typhoon Disaster," *Procedia Computer Science*, vol. 155, pp. 780–784, 2019.
- [21] Q. Guo, Y. Feng, X. Sun, and L. Zhang, "Power Demand Forecasting and Application based on SVR," *Procedia Computer Science*, vol. 122, pp. 269–275, 2017.