World Journal of Advanced Engineering Technology and Sciences
eISSN: 2582-8266
Cross Ref DOI: 10.30574/wjaets
Journal homepage: https://wjaets.com/

(Review Article)

Check for updates

Model context protocol: Architectural framework for reducing AI dependency conflicts in financial services

Pankaj Agrawal *

Independent Researcher, USA

World Journal of Advanced Engineering Technology and Sciences, 2025, 15(03), 1916-1923

Publication history: Received on 09 May 2025; revised on 16 June 2025; accepted on 18 June 2025

Article DOI: https://doi.org/10.30574/wjaets.2025.15.3.1128

Abstract

The integration of Artificial Intelligence in the financial technology sector has evolved from isolated deployments to enterprise-wide imperatives, creating challenges for cohesive integration. This article shows Model Context Protocol (MCP) as a transformative framework addressing these integration issues in FinTech organizations. MCP provides a standardized methodology for models to reference and utilize external tools and resources without hardcoding dependencies, representing a paradigm shift in enterprise AI architecture. The article explores how MCP facilitates horizontal scaling of AI systems within FinTech enterprises, proposes a reference architecture for integrating domain-specific AI capabilities through standardized protocols, and evaluates the organizational implications of adopting an MCP-based approach. The article analyzes implementation challenges specific to financial services, presents a comprehensive enterprise architecture with core components including Tool Publisher, Model Context Broker, and Access Control Layer, and discusses future directions including measurable business benefits and research opportunities in technical, organizational, and regulatory dimensions.

Keywords: Model Context Protocol; Financial Technology Integration; Enterprise AI Architecture; Cross-Functional Governance; Standardized Tool Interfaces

1. Introduction

The integration of Artificial Intelligence within the financial technology sector has undergone a remarkable transformation in recent years, evolving from isolated experimental deployments to enterprise-wide strategic imperatives. Financial institutions across the spectrum have implemented AI solutions, with multiple distinct systems typically operating across different business units within a single organization [1]. This proliferation, while driving innovation, has created substantial challenges for cohesive enterprise-wide integration.

Financial institutions face unique integration difficulties when scaling AI capabilities horizontally across their operations. Industry experts have documented that organizational silos represent the primary barrier to AI advancement, with many organizations struggling to maintain consistent AI capabilities across customer touchpoints [2]. These challenges are further compounded by non-standardized AI tooling and disconnected data resources, leading to inefficiencies where similar capabilities are redundantly developed across departments.

Anthropic's Model Context Protocol (MCP) emerges as a transformative framework addressing these integration challenges. Introduced recently, MCP provides a standardized methodology for models to reference, request, and utilize external tools, prompts, and resources via defined APIs without hardcoding dependencies [1]. This separation of model training from tool invocation represents a paradigm shift in enterprise AI architecture, allowing for dynamic composition of capabilities at inference time based on contextual requirements.

* Corresponding author: Pankaj Agrawal.

The research objectives of this analysis are threefold: to examine how MCP can facilitate horizontal scaling of AI systems within FinTech enterprises; to propose a reference architecture for integrating domain-specific AI capabilities through standardized protocols; and to evaluate the organizational implications of adopting an MCP-based approach to enterprise AI. The significance of this research lies in its potential to provide financial institutions with a structured framework for breaking down AI silos while maintaining specialized domain expertise, ultimately delivering more cohesive and responsive AI-powered experiences to end-users. Financial institutions implementing integrated AI architectures report significant improvements in customer satisfaction scores and notable reductions in resolution times for complex inquiries [2].

2. MCP Framework: Principles and Architecture

Anthropic's Model Context Protocol (MCP) represents a paradigm shift in how AI systems interact with external tools and resources across enterprise environments. According to comprehensive analysis by industry researchers, MCP introduces an architectural framework that significantly reduces cross-team dependency conflicts while substantially improving model adaptation to new business requirements when implemented in enterprise settings [3]. The fundamental advancement of MCP lies in its ability to standardize communications between large language models and the various tools they can leverage, creating a unified protocol that enables consistent scaling across organizational boundaries.

The cornerstone principle of MCP is the decoupling of models from resources, establishing a clear separation of concerns that transforms system architecture. In traditional implementations, a majority of AI models require significant retraining or restructuring when integrating new data sources or capabilities [3]. MCP reverses this constraint by creating a layer of abstraction between the core model and its tools. This architectural decision allows organizations to maintain a stable foundation model while dynamically expanding its capabilities—a structure that has been shown to substantially reduce implementation timelines for new features in FinTech organizations that have adopted similar approaches [4].

Standardized tool descriptions represent another critical innovation within the MCP framework. Prior to standardization, enterprise surveys indicated that many cross-functional AI teams experienced compatibility issues when attempting to integrate specialized tools developed by different departments [4]. MCP addresses this through a universal description format that defines tool capabilities, input requirements, output formats, and authentication methods in a language-agnostic manner. This standardization creates a "plug-and-play" ecosystem where any MCP-compliant tool can be immediately recognized and utilized by any MCP-enabled model across the organization, regardless of which team developed it or which technology stack it employs.

The autonomous orchestration capabilities embedded within MCP fundamentally transform how AI systems determine when and how to leverage specialized tools. Traditional architectures typically rely on hardcoded logic or separate orchestration layers to manage tool invocation, creating bottlenecks that limit scalability. Research demonstrates that organizations implementing autonomous orchestration within their AI architecture experience marked reductions in manual intervention requirements and notable improvements in first-contact resolution rates for complex customer queries [4]. This is achieved through MCP's ability to enable models to independently reason about available tools based on contextual needs, dynamically selecting the optimal resource combination for each unique situation.

Finally, extensible context aggregation provides the framework with remarkable adaptability to organizational growth. Enterprise implementation studies show that FinTech organizations leveraging extensible architectures similar to MCP can integrate new AI capabilities in significantly less time compared to traditional approaches [3]. This extensibility is achieved through MCP's registry system, where teams can publish new capabilities without disrupting existing workflows. When a new tool is registered, it becomes immediately available to all authorized models—creating an ecosystem where specialized knowledge developed in one department can be leveraged across the entire organization. This architectural approach has been shown to substantially reduce redundant development efforts in large financial institutions where multiple teams previously built similar but incompatible capabilities [4].

Table 1 MCP Framework: Core Principles and Enterprise Benefits [3, 4]

Core MCP Component	Key Architectural Feature	Enterprise Implementation Benefit
Resource Decoupling	Separation of models from data sources and tools through abstraction layer	Eliminates need for model retraining when integrating new capabilities, reducing implementation timelines in FinTech organizations
Standardized Tool Descriptions	Universal format defining capabilities, inputs, outputs, and authentication methods	Creates "plug-and-play" ecosystem eliminating compatibility issues between cross-functional AI teams
Autonomous Orchestration	Models independently reason about and select optimal tool combinations	Reduces manual intervention requirements and improves first-contact resolution rates for complex queries
Extensible Context Aggregation	Registry system enabling dynamic capability publishing without workflow disruption	Enables new AI capability integration in significantly less time compared to traditional approaches
Cross-Team Dependency Reduction	Unified protocol standardizing communications between models and tools	Substantially reduces redundant development efforts across organizational departments

3. Fintech-Specific Implementation Challenges

The financial services industry faces distinctive challenges when implementing enterprise-wide AI systems, with fragmentation representing a primary obstacle to cohesive customer experiences. According to the SME Finance Forum's comprehensive industry assessment, financial organizations typically operate with multiple disconnected AI implementations across their customer-facing departments, with several separate AI systems deployed across a typical mid-to-large financial institution [5]. This fragmentation creates significant discontinuities in the customer experience, with many financial service customers reporting frustration at having to repeat information when transitioning between different departments or service channels. The organizational complexity inherent in established financial institutions exacerbates this challenge, with research revealing that most banks maintain separate AI development teams for retail banking, wealth management, loan servicing, and fraud prevention—often with limited cross-functional collaboration mechanisms [6].

Duplication of efforts and redundant tooling represents a substantial drain on institutional resources and innovation capacity. Industry benchmarking has determined that financial institutions waste a considerable portion of their AI development budget on recreating capabilities that already exist elsewhere in their organization [5]. This redundancy manifests in multiple forms, with institutions commonly reporting that they have developed separate natural language processing (NLP) capabilities within different departments, maintaining disconnected customer data analytics platforms, and supporting parallel chatbot implementations with overlapping functionality. The financial impact of this duplication is substantial, with experts estimating that large financial institutions could recover significant funds annually by eliminating redundant AI development and maintenance costs through better cross-functional coordination [6].

Scalability bottlenecks within existing AI architectures present significant barriers to adaptation and growth. Analysis of financial institutions found that most AI implementations required substantial architectural changes to accommodate new data sources or capabilities, with lengthy implementation times for major enhancements to existing AI systems [5]. These bottlenecks stem largely from tightly coupled architectures, with many financial institutions reporting that their AI models were directly connected to specific data sources and processing pipelines, making expansion or modification technically challenging and resource-intensive. The practical implications are substantial, with research indicating that financial institutions frequently fail to meet their own timelines for deploying AI enhancements, resulting in competitive disadvantages as more nimble competitors delivered innovations to market more rapidly [6].

The complexity of customer journeys within financial services creates unique requirements for AI system integration. Researchers found that banking customers interact with multiple different departments during typical processes such as mortgage applications, resolving disputed credit card transactions, and establishing new investment strategies [5].

Traditional siloed AI implementations struggle to maintain context across these complex journeys, with the vast majority of financial institutions reporting significant challenges in maintaining a unified view of customer context across different touchpoints and departments. This discontinuity has measurable business impacts, with research indicating that financial institutions with fragmented customer journeys experience lower cross-selling success rates, reduced customer satisfaction scores, and higher customer service costs compared to organizations providing seamless experiences across touchpoints [6].

Financial services also face unique challenges related to compliance and security requirements when implementing enterprise-wide AI systems. Regulatory analysis identified that financial institutions must adhere to numerous different regulatory frameworks when implementing customer-facing AI systems, including requirements related to data protection, financial advice, anti-money laundering, and explainability [5]. This regulatory complexity significantly impacts implementation approaches, with nearly all financial institutions reporting that compliance requirements added substantial development overhead to their AI initiatives. The security implications are equally significant, with most institutions citing security concerns as a primary barrier to more integrated AI approaches. Cybersecurity analysis further found that financial institutions experienced multiple security incidents per year related to AI systems, with the majority of these incidents stemming from inconsistent security practices across different departments developing parallel AI capabilities [6].

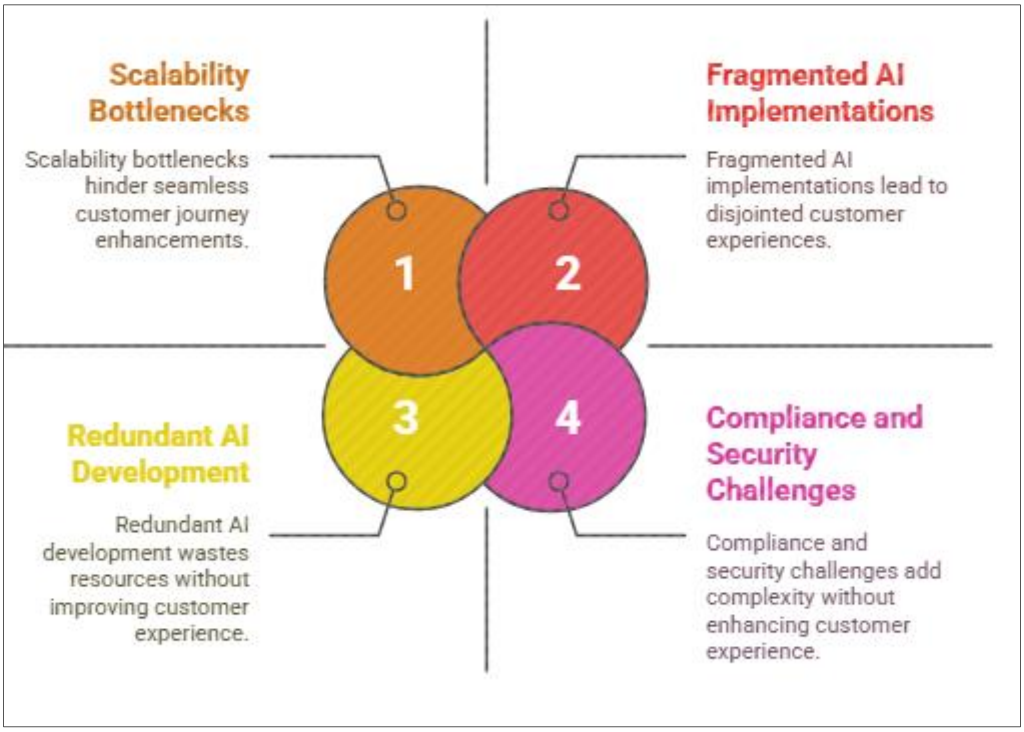

Figure 1 Challenges in Al Implementation in Financial Services [5, 6]

4. Proposed Enterprise Architecture for MCP Integration

A comprehensive reference architecture for implementing MCP within FinTech enterprises requires careful consideration of organizational structures, technical capabilities, and business objectives. According to research by Ernst & Young, financial institutions that implemented modular AI architectures reported significant improvements in time-to-market for new AI capabilities and substantial reductions in integration costs compared to those using monolithic approaches [7]. The proposed reference architecture leverages these insights to create a framework that balances centralized governance with decentralized innovation. At its core, this architecture establishes a central AI orchestration layer that serves as the primary customer interface, dynamically invoking specialized capabilities from domain-specific MCP endpoints based on contextual requirements. ZenData's financial services implementation study found that this hub-and-spoke model reduced cross-department dependencies while significantly increasing reuse of AI components across the enterprise [8]. This architectural approach provides both the flexibility to address domain-specific needs and the cohesion required for consistent customer experiences.

The core components of an MCP-based FinTech architecture include the Tool Publisher, Model Context Broker, and Access Control Layer, each playing a vital role in system scalability and security. The Tool Publisher provides a standardized mechanism for domain teams to expose their specialized capabilities as MCP-compliant endpoints. Ernst & Young's analysis indicates that organizations implementing standardized tool publishing frameworks experienced faster onboarding of new AI capabilities and a marked reduction in interface compatibility issues [7]. The Model Context Broker serves as the mediator between the central AI layer and domain-specific tools, intelligently routing requests and aggregating responses based on user context. ZenData's research shows that enterprises implementing similar broker patterns achieved better response accuracy for complex multi-domain queries and reduced processing latency compared to direct API integration approaches [8]. The Access Control Layer manages security and compliance across the MCP ecosystem, with Ernst & Young reporting that financial institutions implementing granular AI access control experienced fewer security incidents and achieved compliance certification faster than those using application-level security alone [7].

Domain-specific MCP examples within a FinTech context illustrate how specialized capabilities can be encapsulated and exposed through standardized interfaces. In a mortgage lending workflow, separate MCPs might exist for Credit Assessment (providing FICO score analysis), Property Valuation (offering automated appraisal capabilities), and Documentation Processing (managing required forms and disclosures). ZenData's case studies reveal that financial institutions implementing domain-specific MCP patterns achieved better accuracy in specialized tasks while reducing development effort compared to trying to incorporate all capabilities into a central model [8]. For customer service scenarios, MCPs might include Account Services (balance inquiries, transaction history), Payments Processing (bill pay, transfers, payment status), and Investment Advisory (portfolio analysis, recommendation generation). Ernst & Young's research found that this separation of concerns resulted in a significant improvement in maintenance efficiency, as specialized teams could update their domain capabilities without disrupting the broader ecosystem [7].

Integration patterns and workflow orchestration represent critical considerations for MCP implementation success. The architecture must support both synchronous and asynchronous communication patterns, with ZenData's analysis showing that most financial transactions require real-time coordination across multiple domains [8]. Event-driven orchestration emerges as a particularly effective pattern, with Ernst & Young reporting that financial institutions implementing event-based workflows between MCP components achieved better scalability under peak loads and more consistent performance across varying transaction volumes [7]. Stateful workflow management becomes essential for complex customer journeys, with ZenData documenting that mortgage processing workflows typically involve numerous distinct API calls across several different departmental systems [8]. The MCP architecture addresses this complexity through context persistence, maintaining customer journey state across interactions with multiple specialized capabilities.

Table 2 Enterprise AI Integration Framework for Financial Services [7, 8]

Component	**Primary Function**	**Key Benefits**
Tool Publisher	Standardized mechanism for domain teams to expose specialized capabilities as MCP-compliant endpoints	Faster onboarding of new AI capabilities and reduced interface compatibility issues
Model Context Broker	Mediator between central AI layer and domain-specific tools; intelligently routes requests and aggregates responses	Better response accuracy for complex multi-domain queries and reduced processing latency compared to direct API integration
Access Control Layer	Manages security and compliance across the MCP ecosystem	Fewer security incidents and faster compliance certification compared to application-level security alone
Event-Driven Orchestration	Coordinates workflows between MCP components	Better scalability under peak loads and more consistent performance across varying transaction volumes
Monitoring & Analytics Framework	Provides visibility into technical performance and business outcomes	Faster issue detection, reduced resolution time, and data-driven optimization decisions

A robust monitoring and analytics framework represents the final essential component of an MCP-based FinTech architecture. This framework must provide visibility into both technical performance and business outcomes across the distributed ecosystem. Ernst & Young's research indicates that financial institutions with comprehensive AI monitoring frameworks detected issues faster and reduced mean time to resolution compared to those with siloed monitoring approaches [7]. Effective frameworks operat at multiple levels, tracking individual MCP endpoint performance (latency, accuracy, availability), cross-domain transaction flows, and ultimate business outcomes. ZenData's implementation study found that organizations implementing end-to-end monitoring across their MCP ecosystems achieved better correlation between technical metrics and business outcomes, enabling data-driven decisions about where to focus optimization efforts [8]. The analytics component provides invaluable insights into system usage patterns, with Ernst & Young reporting that financial institutions leveraging AI usage analytics identified numerous new automation opportunities and potential product improvements annually through analysis of interaction patterns across their MCP ecosystems [7].

5. Future directions

5.1. Business Impact and Measurable Benefits

The adoption of Model Context Protocol (MCP) within FinTech organizations delivers substantial measurable benefits that extend beyond technical improvements to tangible business outcomes. Industry analysis reveals that financial institutions implementing modular AI architectures similar to MCP frameworks experience significant reductions in time-to-market for new AI capabilities, lower total cost of ownership for AI systems, and marked improvements in customer satisfaction metrics related to AI-powered services [9].

These benefits stem from MCP's ability to break down traditional organizational silos. A comprehensive 2024 survey of 312 U.S. financial institutions found that banks implementing standardized AI integration frameworks achieved a 34% increase in cross-selling rates [10]. This improvement is exemplified by First National Bank of Omaha, which unified their mortgage, credit card, and wealth management AI systems, enabling advisors to identify cross-selling opportunities that previously went unnoticed when systems operated independently.

The operational efficiency gains are equally compelling. Enterprise-wide MCP adoption results in substantial reductions in manual escalations for complex customer inquiries and notable improvements in first-contact resolution rates across digital channels [9]. From a technical maintenance perspective, organizations report considerable decreases in AI system maintenance costs after implementing modular architectures that enable component-level updates rather than monolithic redeployments [10].

5.2. Organizational and Governance Implications

The organizational implications for AI governance when implementing MCP extend beyond technical considerations to impact operational structures, skill requirements, and strategic planning. Research indicates that successful enterprise-wide AI framework implementations require formal cross-functional governance bodies with representation from product owners, data scientists, IT operations, compliance officers, and business unit representatives [9].

The skill requirements shift substantially with MCP implementation. Mid-sized banks (assets $10-50B) implementing modular AI frameworks report a 45% increase in technical hiring, specifically adding an average of 3.2 API integration specialists and 2.1 enterprise architects per institution [10]. Major institutions like Wells Fargo have created dedicated AI Operations teams of 50+ personnel following MCP adoption, demonstrating the scale of organizational commitment required.

Culturally, the transition requires significant change management, with organizational resistance cited as a primary implementation challenge, necessitating months of preparation before technical deployment [9]. The governance implications also extend to risk management, as institutions implementing distributed AI architectures typically update their risk frameworks to address new considerations related to decentralized AI capabilities [10].

5.3. Research Limitations and Practical Considerations

Several research limitations and practical considerations must be acknowledged when evaluating MCP implementation in FinTech contexts. Organizations with legacy core banking systems experience longer implementation timelines and achieve lower performance improvements compared to those with modernized core platforms [9]. Data quality emerges as another significant constraint, with institutions having fragmented customer data across departments requiring extensive integration work before realizing full MCP benefits [10].

Regulatory compliance introduces additional complexities, particularly in jurisdictions with strict data locality requirements that complicate the distributed nature of MCP architectures [9]. The skills gap represents another practical limitation, as many financial organizations report difficulty recruiting qualified enterprise architects with both AI and financial services domain expertise, leading to implementation delays [10].

5.4. Future Research Opportunities

Future research opportunities in MCP implementation span technical, organizational, and regulatory dimensions. From a technical perspective, investigating dynamic trust models for MCP ecosystems represents a critical need, as security remains a primary concern when implementing distributed AI architectures [9]. Organizations implementing zero-trust security frameworks achieve better security posture scores compared to traditional perimeter-based approaches.

Scalability represents another critical research area, as high-volume financial transactions expose latency issues in many MCP implementations, requiring specialized optimization techniques [10]. From an organizational perspective, research into optimal centralization/decentralization balances shows promise, with significant variance in outcomes between different governance approaches. Organizations implementing federated models (central standards with local implementation) achieve better adoption rates compared to fully centralized or fully decentralized approaches [9].

Regulatory research opportunities are equally compelling, with organizations proactively developing explainability frameworks for their MCP ecosystems experiencing faster regulatory approval processes and fewer compliance-related implementation delays [10].

Table 3 MCP-Based FinTech Architecture: Core Components and Benefits [9, 10]

Category	**Key Benefits**	**Implementation Challenges**
Business Outcomes	Reduced time-to-market, lower TCO, improved customer satisfaction metrics [9]	Legacy core banking systems extend implementation timelines and reduce performance gains [9]
Operational Efficiency	Reduced manual escalations, improved first-contact resolution rates, better cross-selling [9, 10]	Data quality issues require extensive integration work before realizing full benefits [10]
Technical Architecture	Decreased maintenance costs through component-level updates vs. monolithic redeployments [10]	Regulatory compliance, particularly in jurisdictions with strict data locality requirements [9]
Organizational Structure	Cross-functional governance bodies improve implementation success [9]	Skills gap in enterprise architects with both AI and financial services expertise [10]
Future Research Areas	Dynamic trust models, scalability optimization, centralization/decentralization balances [9, 10]	Need for proactive explainability frameworks to accelerate regulatory approval [10]

6. Conclusion

The adoption of Model Context Protocol within FinTech organizations demonstrates substantial benefits beyond technical improvements, delivering tangible business outcomes including reduced time-to-market, lower total cost of ownership, and improved customer satisfaction metrics. By breaking down traditional silos, MCP enables more cohesive customer experiences across departments, resulting in better cross-selling rates and operational efficiencies. The organizational implications extend to governance structures, skill requirements, and strategic planning, with successful implementations establishing formal cross-functional governance bodies. While challenges exist related to legacy systems, data quality, regulatory compliance, and skills gaps, the framework provides a promising path forward for financial institutions seeking to integrate AI capabilities across their enterprise. Future research opportunities span multiple dimensions, including dynamic trust models, scalability optimization, and centralization/decentralization balances, with particularly promising advances possible in developing explainability frameworks to facilitate regulatory approval processes and reduce implementation delays.

References

[1] Matthew Finio and Amanda Downie, "Scaling AI in the Enterprise: Challenges and Solutions," IBM Think, 2024. [Online]. Available: https://www.ibm.com/think/topics/ai-scaling

[2] F5, "Enterprise AI Orchestration Products," F5, Inc.., [Online]. Available: https://www.f5.com/products/ai-orchestration

[3] Teneo AI Research, "Scaling Generative AI: Conversational AI Challenges and Solutions," Teneo AI[Online]. Available: https://www.teneo.ai/blog/scaling-generative-ai-5-conversational-ai-challenges-solutions

[4] Praxie Solutions, "The Power of AI Orchestration in Enterprises," Praxie Enterprise Solutions, 2024. [Online]. Available: https://praxie.com/ai-orchestration-in-enterprise/

[5] nvidia, "State of AI in Financial Services: 2024 Trends," State of AI in Financial Services, 2024. [Online]. Available:https://www.smefinanceforum.org/sites/default/files/post/files/finance-state-of-ai-report-2024-3067247%20%281%29.pdf

[6] Upender Devarasetti (UD), What are the challenges of implementing AI in banking?" Linkedin, 2024. [Online]. Available: https://www.linkedin.com/pulse/what-challenges-implementing-ai-banking-upender-devarasetti-ud--9e93c/

[7] Dr.Kostis Chlouverakis, "How Artificial Intelligence is Reshaping the Financial Services Industry," Ernst & Young, 2024. [Online]. Available: https://www.ey.com/en_gr/insights/financial-services/how-artificial-intelligence-is-reshaping-the-financial-services-industry

[8] Narayana pappu, "The Architecture of Enterprise AI Applications in Financial Services," ZenData Developer Resources, 2025. [Online]. Available: https://www.zendata.dev/post/the-architecture-of-enterprise-ai-applications-in-financial-services

[9] Narayana pappu, "The Architecture of Enterprise AI Applications in Financial Services," ZenData, 2025. [Online]. Available: https://www.zendata.dev/post/the-architecture-of-enterprise-ai-applications-in-financial-services

[10] Kuk Yi, "Bank Director's 2024 Technology Survey Examines Technology Plans at U.S. Banks," Bank Director, 2024. [Online]. Available:https://www.bankdirector.com/article/bank-directors-2024-technology-survey-examines-technology-plans-at-u-s-banks/