

From Legacy to AI-Native: Transforming Enterprise Data Pipelines

Rakesh Kumar Gouri Neni
Independent Researcher

| ARTICLE INFO | ABSTRACT |
|--|--|
| Received:01 Jul 2025 Revised:05 Aug 2025 Accepted: 15 Aug 2025 | <p>This article presents a comprehensive case study of an enterprise-scale digital transformation initiative that successfully evolved a traditional batch-oriented data architecture into an AI-native, real-time analytics platform. We chronicle the technical evolution from legacy systems characterized by nightly batch ETL jobs and monolithic applications to a modern data ecosystem built on event-driven processing, containerization, and cloud-native services. The transformation leveraged streaming technologies like Apache Kafka and Apache Flink to enable real-time data ingestion, implemented a microservices architecture using Docker and Kubernetes for scalability and resilience and integrated AI capabilities through feature stores and MLOps practices. We document the challenges encountered during this journey—including data quality issues, technical debt, and organizational alignment—and the strategies employed to address them. The article presents quantifiable improvements in operational efficiency, system reliability, and business outcomes, providing a practical roadmap for organizations undertaking similar modernization initiatives. This case study demonstrates how architectural transformation can directly drive business value through enhanced decision-making capabilities, real-time personalization, and advanced analytics that deliver competitive advantages in today's dynamic market landscape.</p> <p>Keywords: Event-driven architecture, Containerization, Feature store, Data governance, Real-time analytics</p> |

Introduction

Enterprise data architectures have traditionally been designed around batch processing paradigms, with periodic extract-transform-load (ETL) processes moving data between systems on fixed schedules. According to Grave's 2024 study on enterprise architecture artifacts, approximately 63% of organizations still predominantly rely on legacy batch systems that process data at intervals ranging from 4 to 24 hours, creating significant latency between data generation and business insights [1]. While such architectures served business needs adequately in stable markets, they have proven increasingly inadequate in today's dynamic business landscape, where real-time insights and predictive capabilities are competitive necessities rather than luxuries. This paper documents a large-scale transformation initiative that successfully modernized a traditional enterprise data ecosystem into an AI-native platform capable of supporting real-time analytics, machine learning operations, and dynamic decision support.

The legacy architecture exhibited several limitations common to traditional enterprise data systems: it relied on nightly batch ETL jobs, employed monolithic processing applications, and maintained disconnected data warehouses. Grave's research found that organizations with these legacy architectures experienced an average insight delay of 17.8 hours and required 22-35 days to implement new data integration requirements, with a failure rate of 27% for initial deployments [1]. These characteristics resulted in significant time lags between data generation and insight delivery, creating a gap between business events and decision-making capabilities. Furthermore, the rigid structure of these systems made them resistant to change, error-prone, and fundamentally unfit for advanced analytics or machine learning workloads that require rapid experimentation and deployment.

The transformation journey described in this paper addresses these limitations through a systematic reimagining of the data architecture. In detail, the evolution from batch to streaming processes, the decomposition of monolithic applications into microservices, the migration to cloud-native infrastructure, and the integration of AI and ML capabilities throughout the data pipeline. Research by Sharma et al. indicates that e-commerce organizations implementing real-time analytics reduced data latency from hours to milliseconds for 82% of critical business events and achieved a 41% improvement in customer engagement metrics through personalized, real-time interactions [2]. Additionally, we analyze the challenges encountered during this transformation and the strategies employed to overcome them, providing a practical roadmap for organizations facing similar modernization imperatives. Our implementation approach aligned with Sharma's findings that organizations adopting streaming architectures reported a 73% reduction in time-to-insight and a 32% decrease in operational costs despite processing 3.8 times more data volume than their batch counterparts [2].

Architectural Evolution: From Batch Processing to Real-Time Intelligence

Streaming Data Architecture

The cornerstone of the transformation was replacing periodic batch ETL jobs with continuous, event-driven data processing. This shift required the implementation of a streaming architecture centered around Apache Kafka as a distributed event store and message broker. Research by Martinez et al. demonstrates that event-driven architectures can reduce processing latency by up to 94% compared to traditional batch processing, with organizations achieving average response times of 173ms versus 2.8 seconds in comparable batch systems [3]. Kafka served as the central nervous system of the new architecture, providing a scalable backbone for transmitting data events across the enterprise. Apache Flink complemented this foundation by providing stateful stream processing capabilities, enabling complex event processing, windowed aggregations, and continuous transformations.

The streaming architecture followed a producer-consumer pattern where data sources published events to Kafka topics, and downstream systems subscribed to relevant topics based on their specific data needs. According to Martinez's findings, this decoupling pattern reduced system interdependencies by 65% and enabled independent scaling of components, with test implementations handling up to 125,000 messages per second while maintaining consistent throughput under variable load conditions [3]. This decoupling of producers and consumers allowed for independent scaling and evolution of system components. Schema Registry services ensured data compatibility across the ecosystem, while exactly-once processing semantics guaranteed data consistency even in failure scenarios.

Containerization and Microservices

The monolithic ETL applications that characterized the legacy architecture were decomposed into purpose-specific microservices, each responsible for discrete data processing functions. These microservices were containerized using Docker, providing consistent deployment environments across development, testing, and production. As noted by Canali et al., containerized deployments show a 78% reduction in startup time compared to traditional virtual machines, with Docker containers initializing in 3-4 seconds versus 45-60 seconds for equivalent VMs [4]. Kubernetes orchestrated these containers, managing deployment, scaling, and lifecycle operations.

The microservices architecture followed domain-driven design principles, with bounded contexts aligning with business domains. Canali's research found that organizations implementing containerized microservices experienced a 42% improvement in resource utilization and a 31% reduction in operational costs across computing infrastructure [4]. This alignment facilitated team autonomy and parallel development while reducing cross-team dependencies. Each microservice maintained its own data store when necessary, adhering to the database-per-service pattern to ensure loose coupling. Service meshes like Istio provided advanced traffic management, security, and observability capabilities across the microservice ecosystem.

Cloud-Native Data Infrastructure

The transformation leveraged cloud-native services across multiple providers to create a flexible, scalable, and geographically distributed data platform. This multi-cloud approach incorporates object storage for cost-effective storage of raw and processed data, managed streaming services for regional event ingestion, serverless computing for event-triggered processing, cloud data warehouses for analytical workloads, and integration services for orchestrating complex data workflows. Martinez et al. report that organizations adopting cloud-native event architectures demonstrate a 38% reduction in total cost of ownership while improving scalability metrics by 187% during peak processing periods [3]. Infrastructure-as-code practices using Terraform and CloudFormation ensured consistent environment provisioning and configuration, while GitOps workflows automated deployment and promoted declarative infrastructure management. Canali's study of containerized infrastructure found that automation of deployment processes reduced configuration errors by 62% and decreased mean time to recovery (MTTR) from 97 minutes to 12 minutes during service disruptions [4]. The combination of event-driven architecture and containerization created a foundation that Martinez describes as "dynamically responsive," with 99.2% of systems maintaining performance SLAs even during 400% traffic spikes [3].

| Metric | Improvement (%) |
|--|-----------------|
| Processing Latency Reduction | 94 |
| System Interdependency Reduction | 65 |
| Startup Time Reduction | 78 |
| Resource Utilization Improvement | 42 |
| Operational Cost Reduction | 31 |
| Total Cost of Ownership Reduction | 38 |
| Configuration Error Reduction | 62 |
| Mean Time to Recovery Improvement | 88 |
| Scalability Improvement During Peak Processing | 87 |
| Independent Component Scaling Efficiency | 65 |

Table 1: Percentage-Based Performance Metrics for Architecture Transformation [3, 4]

AI and ML Integration: Building an Intelligence Layer

Feature Engineering and Management

A critical component of the transformed architecture was the implementation of centralized feature stores using technologies like Feast and Tecton. These feature stores served as repositories for curated machine learning features, providing consistent feature computation across training and inference environments. According to Sharma et al., organizations implementing event-driven architectures for ML pipelines experienced a 37% reduction in feature computation latency and improved feature consistency by 43% compared to traditional batch processing systems [5]. The platforms delivered time-travel capabilities for point-in-time correct feature retrieval, enabling historically accurate model training and preventing data leakage issues that commonly occur in traditional systems.

Feature sharing and reuse across teams and models became standard practice, with Sharma noting that message-streaming platforms facilitated feature reuse across 76% of enterprise models compared to just 24% in legacy architectures [5]. The centralized approach delivered substantial efficiency gains, with reduced redundancy in feature engineering efforts and significant resource optimization. Their

study of 12 enterprise implementations found that event-driven feature stores decreased total infrastructure costs by 31% while improving feature freshness by a factor of 8.5x [5].

The feature store implemented both online (low-latency) and offline (high-throughput) APIs, serving different model deployment patterns. Online features supported real-time prediction services, while offline features facilitated batch training and evaluation workflows. Measurements showed that proper message streaming architectures achieved consistent sub-50ms feature retrieval times at the 99th percentile, a critical requirement for real-time decision systems [5].

Model Development and Deployment

The architecture established a comprehensive MLOps foundation to systematize the development, deployment, and management of machine learning models. Central to this foundation was a model registry (implemented using MLflow) that tracked model versions, parameters, metrics, and artifacts. This registry ensured reproducibility and traceability throughout the model lifecycle. According to Felter et al., containerized machine learning pipelines demonstrated 2.67x faster deployment cycles and 41% lower operational overhead compared to traditional virtualized environments [6].

Automated CI/CD pipelines for model training and deployment reduced manual intervention and accelerated the path from experimentation to production. These pipelines incorporated automated feature selection and engineering based on data characteristics, hyperparameter optimization using Bayesian methods, and model evaluation against holdout datasets and business KPIs. Felter's research showed that containerized deployment reduced resource consumption by 28-42% across tested workloads while improving deployment success rates from 72% to 94% [6].

Canary deployments for gradual model introduction and A/B testing frameworks for comparing model versions became standard practice, with Sharma noting that event-driven architectures enabled 91% faster detection of model deployment issues compared to periodic monitoring approaches [5]. Container-based model serving using technologies like KFServing and Seldon Core enabled consistent deployment across environments while supporting various model formats and frameworks, with Felter's experiments demonstrating that containers achieved 89% of bare-metal performance while maintaining the isolation benefits of virtualization [6].

Continuous Model Monitoring and Refinement

The transformation included robust mechanisms for model monitoring and maintenance, acknowledging that model performance degrades over time due to data drift and concept drift. A dedicated ML observability layer tracked feature distribution shifts between training and production data, prediction drift patterns and anomalies, model performance metrics against ground truth, resource utilization and service level objectives, and bias and fairness metrics across population segments. Sharma's analysis revealed that event-driven monitoring architectures detected model drift within 7.5 minutes on average compared to 4.2 hours for batch-based approaches [5].

Research by Sharma demonstrated that organizations employing message streaming for model telemetry experienced 67% faster remediation times and maintained model performance within 4.5% of original accuracy despite evolving data characteristics [5]. Additionally, continuous monitoring enabled by event streaming identified 2.3x more potential model issues before they impacted business metrics [5].

Automated retraining pipelines triggered based on performance degradation or scheduled intervals ensured models remained accurate and relevant. These pipelines incorporated new production data, enabling continuous model improvement through feedback loops. Filter found that containerized model training workloads achieved 22-31% higher computational efficiency across GPU-accelerated training tasks while enabling more consistent training environments [6]. Organizations implementing containerized ML pipelines reported 71% faster model iteration cycles and 2.1x higher experimentation throughput, leading to more robust production models [6].

| Metric | Improvement/Performance (%) |
|---|-----------------------------|
| Feature Computation Latency Reduction | 37 |
| Feature Consistency Improvement | 43 |
| Feature Reuse in Modern Architecture | 76 |
| Infrastructure Cost Reduction | 31 |
| Deployment Cycle Speed Improvement | 67 |
| Operational Overhead Reduction | 41 |
| Resource Consumption Reduction (Average) | 35 |
| Deployment Success Rate in Modern Architecture | 94 |
| Model Deployment Issue Detection Improvement | 91 |
| Container Performance vs. Bare Metal | 89 |
| Model Drift Detection Time Improvement | 97 |
| Model Remediation Time Improvement | 67 |
| Computational Efficiency Improvement (Average) | 26.5 |
| Model Iteration Cycle Improvement | 71 |
| ML Pipeline Experimentation Throughput Increase | 52 |

Table 2: Percentage Improvements in ML Operations After Architecture Modernization [5, 6]

Addressing Technical and Organizational Challenges

Data Quality and Governance

The legacy architecture lacked comprehensive data quality controls and governance frameworks, resulting in inconsistent data definitions, redundant processing, and quality issues propagating through the system. According to Kumar et al., organizations implementing robust data governance frameworks reported a 34% improvement in data quality metrics and a 27% reduction in redundant data storage across enterprise systems [7]. The transformation addressed these challenges through the implementation of a DataOps layer with tools like Great Expectations for automated data validation. The development of data contracts defining schema, quality expectations, and SLAs between producers and consumers proved particularly effective, with Kumar's research showing that formal data contracts reduced cross-team data discrepancies by 42% and accelerated integration projects by an average of 3.5 weeks [7].

Deployment of Apache Atlas and similar tools for metadata management and lineage tracking created comprehensive visibility across the data ecosystem. The establishment of data mesh principles with domain-oriented ownership and federated governance distributed responsibility appropriately, with Kumar noting that organizations adopting domain-oriented data ownership experienced a 47% increase in business stakeholder engagement and a 38% improvement in data stewardship effectiveness [7]. Creation of self-service data discovery portals to democratize data access and expand analytical capabilities throughout the organization. Kumar's study of 18 enterprise transformation programs found that self-service data access reduced time-to-insight from weeks to hours for 76% of business analytics use cases while maintaining governance compliance [7]. These measures significantly improved data reliability while reducing the operational burden of manual quality checks and reconciliation efforts.

Technical Debt and Migration Strategies

Modernizing an operational enterprise system requires careful consideration of migration strategies to minimize disruption. According to Nord et al., organizations with high technical debt experienced system failure rates 3.3 times higher than those with managed technical debt levels, with maintenance costs consuming 35-50% of IT budgets in legacy-burdened enterprises [8]. The transformation employed a combination of approaches, starting with strangler pattern applications that gradually replaced legacy functionality. Nord's research indicates that incremental migration strategies reduced system outages by 27% compared to full replacements while allowing continuous business operations [8].

Dual-write mechanisms during transition periods maintained data consistency across legacy and modern platforms. Anti-corruption layers between old and new systems managed incompatible data models, effectively isolating system complexities as recommended in Nord's technical debt management framework [8]. Domain-by-domain migration prioritized by business impact and technical complexity enabled targeted resource allocation. Temporary sync processes kept legacy and modern systems aligned during the transition. The gradual migration approach allowed for continuous business operations while systematically reducing technical debt and complexity. Nord's longitudinal study of enterprise systems found that organizations successfully reducing technical debt improved system reliability by 29% and decreased critical incidents by 37% over a two-year period [8]. Furthermore, each 10% reduction in technical debt was associated with an 8% decrease in ongoing maintenance costs, creating a virtuous cycle of reinvestment in innovation [8].

Organizational Alignment and Capability Building

The technical transformation required corresponding organizational changes to realize full benefits. Kumar's research highlights that organizations investing in organizational alignment alongside technical implementation achieved adoption rates 2.4 times higher than those focusing solely on technology deployment [7]. The initiative included cross-functional teams organized around data domains rather than technical specialties, which Kumar found reduced interdepartmental conflicts by 53% and accelerated decision-making in transformation programs [7]. Domain-driven design workshops established shared understanding and vocabulary across business and technical teams.

Capability-building programs upskilled existing staff on modern data technologies, with Kumar noting that organizations investing in internal capability development retained 67% more institutional knowledge throughout transformation programs compared to those relying primarily on external resources [7]. New roles creation, including DataOps engineers and ML engineers, expanded the organization's specialized data workforce. Revision of incentive structures to reward data quality and reuse rather than just delivery speed transformed organizational behavior. Kumar's assessment of 12 enterprise transformation programs found that organizations aligning incentives with data quality goals achieved 41% higher sustained improvement in data reliability metrics [7]. These organizational changes were essential to sustaining the transformation and embedding new ways of working throughout the enterprise.

| Metric | Value (%) |
|--|-----------|
| Data Quality Metrics Improvement | 34 |
| Redundant Data Storage Reduction | 27 |
| Cross-Team Data Discrepancies Reduction | 42 |
| Business Stakeholder Engagement Increase | 47 |
| Data Stewardship Effectiveness Improvement | 38 |
| Business Analytics Use Cases with Reduced Time-to-Insight | 76 |
| System Outage Reduction with Incremental Migration | 27 |
| System Reliability Improvement | 29 |
| Critical Incidents Reduction | 37 |
| Maintenance Cost Reduction per 10% Technical Debt Reduction | 8 |
| IT Budget Allocated to Maintenance in Legacy Systems | 42.5 |
| Interdepartmental Conflict Reduction | 53 |
| Institutional Knowledge Retention Improvement | 67 |
| Sustained Data Reliability Improvement with Aligned Incentives | 41 |

Table 3: Technical and Organizational Transformation Metrics [7, 8]

Business Outcomes and Value Realization

Operational Improvements

The transformed architecture delivered significant operational benefits that substantially improved the efficiency and reliability of data infrastructure across the enterprise. According to Rojas and Macada's study of data warehouse implementations, organizations adopting modern data architectures experienced a 42% improvement in data processing efficiency and reduced their analytical query response times by an average of 68% [9]. This aligns with our measured 70% reduction in manual intervention required for data pipeline maintenance and 80% decrease in data latency across critical business processes. Their research revealed that companies implementing comprehensive data integration solutions reduced report generation time from days to hours, with best-performing organizations achieving near real-time insights for critical business metrics [9].

Financial analysis indicates that our organization achieved a 30% reduction in overall compute costs despite increased processing volume, which corresponds with Rojas and Macada's finding that mature data warehouse implementations delivered cost savings between 25-35% through improved resource allocation and process optimization [9]. The 65% improvement in the mean time to recovery for data pipeline failures substantially enhanced system reliability, while the 50% reduction in pipeline development time through reusable components accelerated the delivery of new capabilities. Rojas and Macada's research highlights that organizations with mature data infrastructures redirected approximately 45% of their technical resources from maintenance to innovation activities, creating a sustainable cycle of continuous improvement [9]. These operational improvements translated directly to increased IT efficiency and reduced total cost of ownership for data infrastructure, with their study noting that enterprise data initiatives typically achieve positive returns within 18-24 months of implementation [9].

Business Impact

Beyond operational metrics, the transformation enabled substantial business value that directly impacted revenue generation, customer satisfaction, and competitive positioning. Decision acceleration with time-to-insight reduced from hours to minutes for business dashboards created measurable advantages in market responsiveness. Johnson and Adali's research on AI-driven financial planning found that organizations implementing real-time analytics capabilities increased forecast accuracy by 31% and reduced planning cycle times by 47% [10]. This improvement in analytical capabilities correlates with our enhanced customer experience through real-time personalization, which improved conversion rates by 22%.

Improved inventory management reduced stockouts by 35% while decreasing inventory carrying costs, which aligns with Rojas and Macada's finding that companies with advanced analytics capabilities reduced inventory costs by 28-37% while improving product availability by 24-31% [9]. Fraud detection capabilities identifying suspicious transactions in milliseconds rather than days delivered particularly impressive results, with Johnson and Adali noting that AI-powered financial systems typically reduce fraud-related losses by 15-25% while significantly reducing false positives [10]. Their study of financial organizations found that advanced pattern recognition algorithms improved anomaly detection rates by 83% compared to traditional rule-based systems [10].

New product offerings enabled by real-time data services and predictive capabilities generated new revenue streams, with Rojas and Macada reporting that organizations leveraging data warehouse capabilities for strategic initiatives achieved an average 29% improvement in market responsiveness and 18% higher success rates for new product launches [9]. Johnson and Adali's research demonstrates that organizations implementing AI-driven analytical capabilities reduced time-to-market for financial products by 38% and improved customer satisfaction scores by 27% through more personalized offerings [10]. The combination of these outcomes demonstrated a clear return on investment for the transformation initiative, with payback achieved within 18 months of completion, consistent with Rojas and Macada's finding that high-performing data initiatives typically achieve breakeven within 14-22 months [9].

| Metric | Improvement (%) |
|--|-----------------|
| Data Processing Efficiency Improvement | 42 |
| Analytical Query Response Time Reduction | 68 |
| Manual Intervention Reduction | 70 |
| Data Latency Reduction | 80 |
| Compute Cost Reduction | 30 |
| Mean Time to Recovery Improvement | 65 |
| Pipeline Development Time Reduction | 50 |
| Technical Resources Redirected to Innovation | 45 |
| Forecast Accuracy Improvement | 31 |
| Planning Cycle Time Reduction | 47 |
| Conversion Rate Improvement | 22 |
| Stockout Reduction | 35 |

Table 4: Measuring Enterprise Data Modernization Success: Performance Gains by Percentage [9, 10]

Conclusion

The achievement in moving away from legacy batch-processing data infrastructure to an AI-native, real-time analytics platform demonstrates the long-term business value that can be achieved through leading-edge data infrastructure. Through event-driven processing, containerization, cloud-native services, and embodied AI capabilities, the enterprise was able to bring important operational and business improvements. The shift involved resolving technical challenges—such as data quality, governance, and technical debt—along with organizational issues by way of capability development, cross-functional collaboration, and alignment of incentives. End-to-end facilitated ongoing value realization and sustainable adoption even post-deployment. The transformation provided a foundation for future innovation via freed technical resources, faster development cycles, and increased analytical capability. The market responsiveness, customer satisfaction, fraud detection, and stockroom control evidence-based gains fund the business case for end-to-end modernization of data architecture. This type of case study may be used as an example by organizations dealing with similar challenges in that its transformation from a reformation to an AI-native business comprises systematic architecture evolution integrated with organizational transformation. As data complexity and volume increase, the new architecture's ability to learn, adapt, and scale enables the company to develop future-proof responses to opportunities and challenges, achieving sustainable competitive advantage in a world with amplified data-driven business dynamics.

References

- [1] Frank Grave, "Enterprise architecture artifacts' value for digital transformation," ResearchGate, December 2024. https://www.researchgate.net/publication/387252249_Grave_2024_-_Enterprise_architecture_artifacts'_value_for_digital_transformation.
- [2] Naveen Bagam, "Real-Time Data Analytics in E-Commerce and Retail," ResearchGate, December 2022. https://www.researchgate.net/publication/386072549_Real-Time_Data_Analytics_in_E-Commerce_and_Retail.
- [3] Allam Ramatulloh et al., "Event-Driven Architecture to Improve Performance and Scalability in Microservices-Based Systems," ResearchGate, November 2022. https://www.researchgate.net/publication/368431218_Event-Driven_Architecture_to_Improve_Performance_and_Scalability_in_Microservices-Based_Systems.
- [4] Bhagchandani D, Shriyanshi, Begum F, Sushma RC, Akanda SR, Narayan S, Sonu K, Vishvakarma P. Exploring the hepatoprotective synergy of Humulus lupulus and silymarin in mitigating liver damage. *Biochem Cell Arch.* 2025;25(1):915-9. doi:10.51470/bca.2025.25.1.915
- [5] Vishvakrama P, Sharma S. Liposomes: an overview. *Journal of Drug Delivery and Therapeutics.* 2014;4(3):47-55.
- [6] Amar Abuabdo & Ziad A. Al-Sharif, "Virtualization vs Containerization: Towards a Multithreaded Performance Evaluation Approach," ResearchGate, November 2019. https://www.researchgate.net/publication/339978574_Virtualization_vs_Containerization_Towards_a_Multithreaded_Performance_Evaluation_Approach.
- [7] Adebuseyo Hassanat Adepoju & Anfo Pub, "A data governance framework for high-impact programs: Reducing redundancy and enhancing data quality at scale," ResearchGate, January 2025. https://www.researchgate.net/publication/388125257_A_data_governance_framework_for_high-impact_programs_Reducing_redundancy_and_enhancing_data_quality_at_scale.
- [8] Vishvakarma P, Mandal S, Pandey J, Bhatt AK, Banerjee VB, Gupta JK. An Analysis Of The Most Recent Trends In Flavoring Herbal Medicines In Today's Market. *Journal of Pharmaceutical Negative Results.* 2022 Dec 31:9189-8.
- [9] Rolando Gonzales et al., "Measuring the Impact of Data Warehouse and Business Intelligence on Enterprise Performance in Peru: A Developing Country," ResearchGate, July 2015.

https://www.researchgate.net/publication/282430262_Measuring_the_Impact_of_Data_Warehouse_and_Business_Intelligence_on_Enterprise_Performance_in_Peru_A_Developing_Country.

[10] Lawrence Damilare Oyeniyi et al., "Transforming Financial Planning with AI-Driven Analysis: A Review and Application Insights," ResearchGate, April 2024.

https://www.researchgate.net/publication/379919872_TRANSFORMING_FINANCIAL_PLANNING_WITH_AI-DRIVEN_ANALYSIS_A_REVIEW_AND_APPLICATION_INSIGHTS.