International Journal of Computer Applications (0975 – 8887)
Volume 187 – No.136, August 2026

An MCP-Driven Multi-Agent Enterprise Marketplace Platform Architecture for Myanmar Secondhand and Local Fashion Merchants: A Survey-based Requirements Analysis

Nei Rin Zara Lwin
Department of Computer Science
University of the People
Pasadena, California, USA

Kyaw Kyaw Oo
School of Engineering Practice
McMaster University
Hamilton, Ontario, Canada

Si Thu Aung
Department of Mathematics
State Univ. of New York at Buffalo
Buffalo, New York, USA

ABSTRACT
Myanmar's emerging secondhand and local fashion merchants operate predominantly through Facebook and Viber, lacking automated tools for order management, AI-assisted discovery, or consumer trust verification. This paper presents a survey-based requirements analysis and a corresponding Model Context Protocol (MCP)-driven multi-agent enterprise architecture for a dedicated marketplace platform targeting this segment. A structured needs-assessment survey of 15 merchants identified operational pain points and ranked 16 platform capabilities on a five-point Likert scale. Features scoring 3.85 or above on the survey were adopted as active requirements, producing eleven survey-qualified capabilities. A twelfth — AI condition grading — was added independent of its survey score, since it serves as a trust signal rather than a popularity measure. Together, these twelve capabilities map across five MCP tool servers: analytics, order-manager, recommendation, chat-assistant, and product-listing. The proposed five-layer architecture spans a React Native/Next.js client layer, a LangGraph multi-agent orchestration layer coordinated via the Agent-to-Agent (A2A) protocol, the MCP tool layer, a MongoDB/Redis data infrastructure layer, and an observability and monitoring layer. Trust is the biggest factor driving purchases in Myanmar social commerce. To address this, the architecture builds a trust layer — based on signaling theory — into its orchestration and MCP tool layers, using AI condition grading, seller reputation scores, and verified-seller badges. This approach offers a repeatable model for building AI-first marketplaces in similar emerging markets.

General Terms
Digital Commerce, Artificial Intelligence, Platform Architecture, Developing Economies, MCP, Consumer Trust.

Keywords
Myanmar, secondhand fashion, Model Context Protocol, multi-agent systems, digital transformation, needs assessment, platform architecture, consumer trust, requirements engineering.

1. INTRODUCTION
Myanmar's retail landscape exhibits a pronounced digital asymmetry. Thrift merchants and local fashion brands have moved product discovery onto Facebook Pages and Instagram, yet the operational infrastructure beneath remains informal: orders managed through Messenger inboxes, payments via manual bank transfers, and inventory tracked by human memory. This is not a convenience gap — it is a structural ceiling on merchant growth, revenue predictability, and buyer trust.

This reliance on social media is contextually grounded. As of January 2024, Myanmar had 18.5 million social media users — 33.8% of its population — with 64.28 million active mobile connections [1]. Facebook dominates the commerce landscape by default, not by design. Htet et al. (2024) found that trust is the dominant determinant of consumer buying behavior on Facebook Marketplace in Myanmar, outweighing visual appeal, ease of use, and electronic word-of-mouth [2]. Yet Facebook and Telegram provide no formal trust infrastructure: no verified seller credentials, no standardized condition grades, no product provenance. The result is a market where trust is scarce, and transaction friction is structurally high.

This paper proposes a resolution centered on the Model Context Protocol (MCP) [3] — Anthropic's open JSON-RPC standard enabling AI agents to invoke backend tools in a structured, composable manner. By grounding each architectural decision in empirically identified merchant and consumer needs, the proposed system addresses three compounding deficits simultaneously: operational inefficiency, AI capability gaps, and the trust infrastructure vacuum that limits Myanmar's social commerce ecosystem.

The paper follows a strict evidence chain: (1) a structured merchant needs-assessment survey identifies and ranks operational pain points and desired capabilities; (2) survey findings serve as the primary basis for prioritizing which MCP tool servers are designed and at what priority; (3) trust signal requirements, grounded in signaling theory [4], are embedded as first-class platform architecture components alongside AI automation features.

This paper makes four primary contributions:

1. A structured needs-assessment framework identifying operational requirements of Myanmar secondhand and local fashion merchants.

2. A requirements-to-MCP mapping methodology linking empirical merchant pain points to dedicated intelligent tool server designs.

3. An MCP-driven multi-agent enterprise architecture with development priority determined by survey evidence.

4. A trust-aware architectural design integrating AI automation with consumer trust mechanisms identified in Myanmar social commerce research.

2. BACKGROUND AND RELATED WORK

2.1 Marketplace Evolution and Generation-Four Platforms

For the purposes of this paper, marketplace evolution is conceptualized into four generations: (i) transaction-centric directories (Craigslist, OLX); (ii) feedback-enriched marketplaces (eBay, Carousell); (iii) recommendation-augmented platforms (Poshmark, Mercari); and (iv) emerging intelligence-centric platforms where AI agents perform listing, pricing, and matching autonomously on behalf of sellers. Current market leaders remain largely in generation two or three — machine learning is applied to recommendation engines while the seller's core operational workflow stays manual [5]. Bae et al. (2022) surveyed technology-based strategies on secondhand platforms including Karrot, OfferUp, and Depop, finding that platform-level AI features — not merely product assortment — are the primary differentiators driving adoption [6]. This paper proposes a generation-four architecture for the Myanmar context.

2.2 Myanmar's Digital Commerce Landscape

Myanmar presents a distinctive digital commerce context. Internet penetration stood at 44.0% in January 2024, with 24.11 million internet users and 64.28 million active mobile connections [1]. The digital economy remains characterized by mobile-first access, limited fixed broadband (6.6% household penetration), and a commercial landscape dominated by Facebook and Telegram as de facto storefronts for small merchants [7]. The Economic Research Institute for ASEAN and East Asia (ERIA) identifies ICT infrastructure gaps, low digital literacy, and absence of formal digital payment rails as the primary structural barriers to SME e-commerce adoption in Myanmar [7].

Kyi (2024) confirmed that customer service quality was the strongest predictor and merchandise attributes were also positively associated of satisfaction among Myanmar online shoppers (n = 340) [8], underscoring that the gap between social-media convenience and professional e-commerce standards is felt acutely by buyers. These findings together define the migration problem: merchants need a purpose-built platform that eliminates more friction than the migration creates.

2.3 Consumer Trust in Myanmar Social Commerce

Htet et al. (2024) conducted a quantitative study of 384 Myanmar Facebook Marketplace users and found that trust was the strongest factor influencing buying behavior, outperforming visual appeal, ease of use, and e-WOM [2]. This finding is consistent with broader social commerce literature: Lita et al. (2024) and related ASEAN studies show that trust in the seller mediates the relationship between social commerce constructs and purchase intention across developing-market platforms [9].

These findings inform the proposed architecture: condition grading, seller reputation scoring, and product provenance are modeled as first-class requirements alongside AI automation and operational features.

2.4 Signaling Theory and Secondhand Fashion Trust

Negash and Akhbar (2024) applied signaling theory to 203 consumers on a secondhand fashion platform and found that seller reputation is the strongest trust signal, followed by product history and refurbishment details [4]. Functionality-oriented consumers respond most to remarketing information; newness-conscious consumers respond more to condition disclosures. Environmental awareness significantly moderates the trust-to-engagement relationship — a finding with growing relevance as Myanmar's youth consumer segment shows rising environmental consciousness.

This research grounds two specific architecture decisions: AI-generated condition grading is a trust signal in the sense of signaling theory [4], and seller reputation scoring must be treated as an architecture-level first-class requirement rather than an enhancement.

2.5 Model Context Protocol (MCP)

MCP [3] defines a client-server protocol in which a host LLM connects to MCP servers exposing tools via standardized JSON-RPC. Each tool is described by a machine-readable schema; the AI agent autonomously selects and invokes tools at inference time without hard-coded integration logic. Krishnan (2025) demonstrated that MCP-mediated multi-agent systems achieve improved modularity and auditability over monolithic LLM integrations [10]. Enterprise adoption studies suggest that MCP can reduce integration surface area and accelerate AI capability deployment in production systems [11].

2.6 Multi-Agent Systems for Commerce

Wooldridge's multi-agent framework [12] establishes that complex environments benefit from agent specialization: dedicated agents for listing, pricing, fraud detection, and analytics each carry bounded responsibility, enabling independent evaluation and replacement. The Agent-to-Agent (A2A) protocol [13] provides a standard messaging layer for inter-agent coordination. Multi-agent commerce systems have been shown to reduce mean listing time, improve pricing accuracy, and increase buyer-seller match quality in platform deployments [5].

Existing studies focus on user behavior, technology adoption, or individual platform features. Few studies integrate empirical merchant requirements with an MCP-driven multi-agent enterprise architecture tailored to Myanmar’s secondhand and local fashion ecosystem. This gap motivates the proposed architecture.

3. RESEARCH METHODOLOGY

This study employs a Needs Assessment methodology [14] rather than technology-acceptance models (TAM, UTAUT), because the proposed platform does not yet exist for respondents — speculative acceptance questions would yield methodologically invalid responses. The analytical sequence is:

- Phase 1 — Industry Problem Framing (Myanmar digital commerce context)
- Phase 2 — Merchant Survey (Needs Assessment, approximately twenty merchants)
- Phase 3 — Requirement Analysis (mean scores → ranked list)
- Phase 4 — Architecture Proposal (MCP + multi-agent + trust layer)

The survey instrument comprises 25 fixed-choice questions across three sections: Section A (seller profile), Section B (operational challenge severity, five-point Likert scale (a standard 1–5 rating scale) ranging from 1 (not interested) to 5 (very interested). Human-readable language with no technical jargon is used throughout. Trust infrastructure requirements are embedded in the architecture independently of survey scores, grounded in Sections 2.3–2.4.

4. SURVEY RESULTS

4.1 Seller Profile

Four profile items establish respondent context: selling tenure, product category, current platforms, and monthly order volume. Fifteen merchants participated. Platform dependence is pronounced: thirteen of fifteen respondents sell exclusively via Facebook, confirming that Facebook-native sellers are the highest-priority migration target for the proposed platform [2]. The cohort spans early-stage sellers (under six months, four respondents) through experienced merchants (over three years, six respondents), with product categories concentrated in secondhand and thrift fashion (ten respondents), vintage goods (five), and bags, footwear, and accessories (four). Monthly order volumes range from fewer than ten to over one hundred, reflecting the heterogeneous scale of Myanmar's micro-merchant segment.

4.2 Operational Pain Points

Five operational challenge items were rated on a five-point severity scale, ranging from 1 (no problem) to 5 (a serious daily problem). All five challenges scored at or below the scale midpoint (mean range: 2.40–3.00), indicating that merchants have developed functional workarounds rather than experiencing acute daily crises. The highest-rated challenge was timely response to buyer messages (mean 3.00), reflecting the volume of simultaneous conversations characteristic of Facebook Commerce. Pricing uncertainty followed (mean 2.73), and inventory sell-through visibility (mean 2.67). Order tracking across chat threads (mean 2.40) and product photo quality (mean 2.47) were rated lowest. The moderate pain scores indicate adaptation cost rather than an absence of need — merchants who have built manual workarounds represent a migration target, not an absence of demand.

4.3 Feature Interest

Sixteen feature interest items were rated on a five-point desirability scale (1 = not interested, 5 = very interested), split into basic platform capabilities (five items) and AI-powered capabilities (eleven items). Among basic features, multi-store account management led at mean 4.67, followed by in-app payment integration (4.53) — both pointing to structural friction: merchants managing multiple social pages and performing manual payment slip verification daily. Among AI-powered features, weekly sales intelligence summaries scored highest across all sixteen items at mean 4.73, closely followed by AI buyer-seller matching (4.67) and per-item profit tracking (4.60). These results position analytics and intelligent matching as primary merchant priorities. The AI negotiation reply assistant scored lowest across all features at mean 2.87 — a significant gap from the next-lowest item — consistent with the trust-relationship norms of Myanmar social commerce, where price negotiation carries a personal dimension that merchants are unwilling to delegate to automation. Table 1 presents all sixteen items ranked by mean score. Mean score ranges per MCP server, and their exposed tools are summarized in Table 2.

Scores cluster into two bands separated by a natural break at the 3.85 threshold: eleven items meet or exceed the cut-off and are mapped to five MCP servers; a twelfth item, AI condition grading, is retained as a trust-signal requirement despite falling below the threshold; the remaining four items fall below the threshold and are excluded from the active architecture (Figure 1).

Table 1. Feature Interest Rankings (n = 15)

Feature	Mean	MCP Server
Weekly AI sales summary	4.73	analytics
Multi-store account management	4.67	order-manager
AI buyer-seller matching	4.67	recommendation
Per-item profit tracking	4.60	analytics
In-app payment integration	4.53	order-manager
AI FAQ auto-reply	4.47	chat-assistant
AI sales probability prediction	4.27	recommendation
Unified order dashboard	4.20	order-manager
In-app buyer messaging	4.20	chat-assistant
AI photo enhancement	4.00	product-listing
AI listing generation from photos	3.87	product-listing
AI condition grading	3.53	product-listing
Duplicate listing detection	3.60	excluded
AI price suggestion	3.53	excluded
Make-an-offer workflow	3.47	excluded
AI negotiation assistant	2.87	excluded

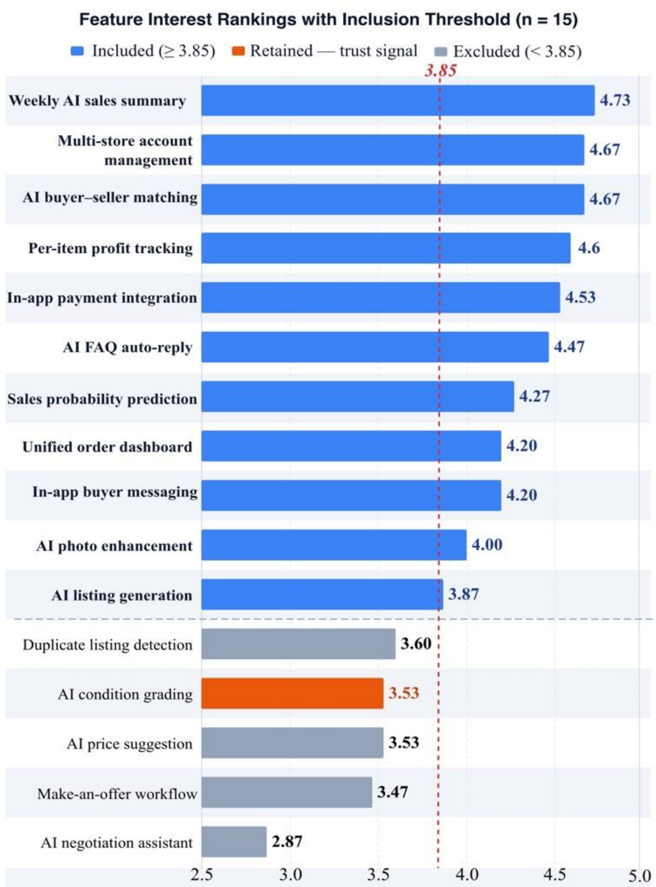

Figure 1. Feature Interest Rankings with Inclusion

The analytics and order-manager servers attract the highest aggregate demand — four of the top six items belong to these two servers — confirming that profit visibility and multi-shop order management are the dominant merchant priorities.

Excluded items share a common profile: pricing and negotiation tasks that merchants currently handle through direct Messenger conversations, flagging this conversational layer as a natural scope for a subsequent release.

5. REQUIREMENTS ANALYSIS

Merchant-reported mean scores determine MCP server development priority, reflected in Table 2. Features with mean scores of 3.85 or above are included as active architectural requirements. The threshold of 3.85 was selected at the natural cluster break in the ranked distribution: a 0.27-point gap separates the eleventh-ranked feature — AI listing generation (3.87) — from the twelfth — duplicate listing detection (3.60) — marking a clear discontinuity between features with strong merchant interest and those with moderate or low interest (Figure 2). This approach follows the discrepancy-boundary principle of Needs Assessment methodology [14], which prescribes placing the cutoff at a meaningful gap in the ranked distribution rather than at an arbitrary scale value. Four features fell below this threshold — AI offer negotiation (2.87), make-an-offer workflow (3.47), AI price suggestion (3.53), and duplicate listing detection (3.60) — and are excluded from the initial architecture scope. AI condition grading (mean 3.53) is retained as an architectural requirement independent of its survey score, as it serves a direct trust-signal function grounded in signaling theory [4]. Trust requirements TR-1 through TR-3 are embedded as architecture-level requirements independent of survey scores, grounded in Sections 2.3–2.4.

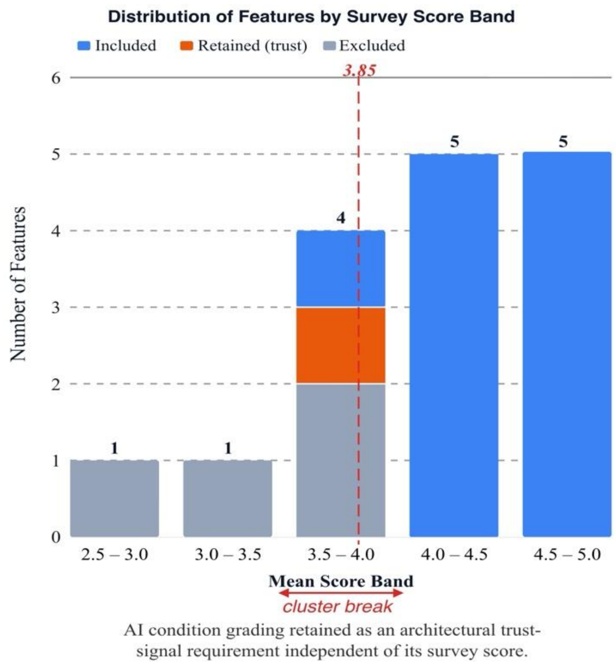

Figure 2. Distribution of Features by Survey Score Band

The twelve included features group naturally into five bounded domains, each mapped to a dedicated MCP server. The analytics server aggregates the two highest-demand capabilities — weekly AI sales summary (4.73) and per-item profit tracking (4.60) — reflecting merchants' strongest expressed need for business intelligence. The order-manager server consolidates three operational workflow features: multi-store account management (4.67), in-app payment integration (4.53), and unified order dashboard (4.20), addressing the manual coordination burden that currently spans Messenger, banking apps, and spreadsheets. The recommendation server pairs AI buyer-seller matching (4.67) with sales probability prediction (4.27), both serving inventory intelligence: the first surfaces the right product to the right buyer, the second tells the merchant which stock to prioritize. The chat-assistant server handles AI FAQ auto-reply (4.47) and in-app buyer messaging (4.20), replacing Messenger as the primary communication channel. Finally, the product-listing server covers AI photo enhancement (4.00) and AI listing generation (3.87) — the two features that reduce listing friction — alongside AI condition grading (3.53), retained as a trust-signal requirement independent of its survey score. This grouping ensures each MCP server has a coherent, bounded responsibility and that development priority follows directly from descending survey evidence, as summarized in Figure 3.

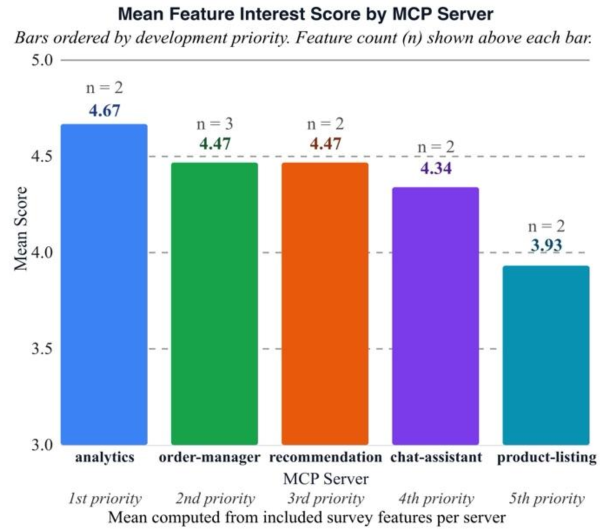

Figure 3. Mean Feature Interest Score by MCP Server, ordered by development priority.

6. PROPOSED ARCHITECTURE

The proposed architecture is organized into five horizontal layers. The Client Layer serves React Native mobile and Next.js admin and public frontends, communicating with the orchestration layer only through a versioned REST and WebSocket API to keep client concerns separate from agent internals. The Agent Orchestration Layer, built in LangGraph [15] on top of the LangChain framework [16], manages stateful, cyclical agent workflows and coordinates the Seller, Buyer, and Order agent families via the Google A2A Protocol [13], allowing them to collaborate without direct coupling. The MCP Tool Layer comprises five domain-specific MCP servers that expose backend capabilities as discrete, versioned JSON-RPC tools through a central registry [3]. The Data and Infrastructure Layer combines MongoDB for document storage, Redis for session and cache management, BullMQ for asynchronous AI inference jobs, and Cloudflare R2 for product images — each matched to the access pattern of the workload it serves. The Observability and Monitoring Layer rounds out the stack: Prometheus and Grafana handle metrics and alerting, OpenTelemetry and Jaeger provide distributed tracing, Loki aggregates logs, and LangFuse [18] adds LLM-specific observability — prompt traces, token usage, and cost — across every agent and MCP tool call, while MLflow [19] tracks the experiment runs and versioned models behind the recommendation service.

A Trust Layer cuts across the orchestration and MCP tool layers rather than standing as a sixth layer. The Governance agent family tracks order completion rates and response times to keep reputation signals current, while the product-listing server exposes gradeCondition as a persistent trust artifact on each product record. Computing trust signals in the agent layer and storing trust artifacts in the MCP layer keeps buyer-facing signals consistent regardless of which client or agent initiates the query [4]. Figure 4 shows all five layers together, with this trust layer embedded across layers 2 and 3.

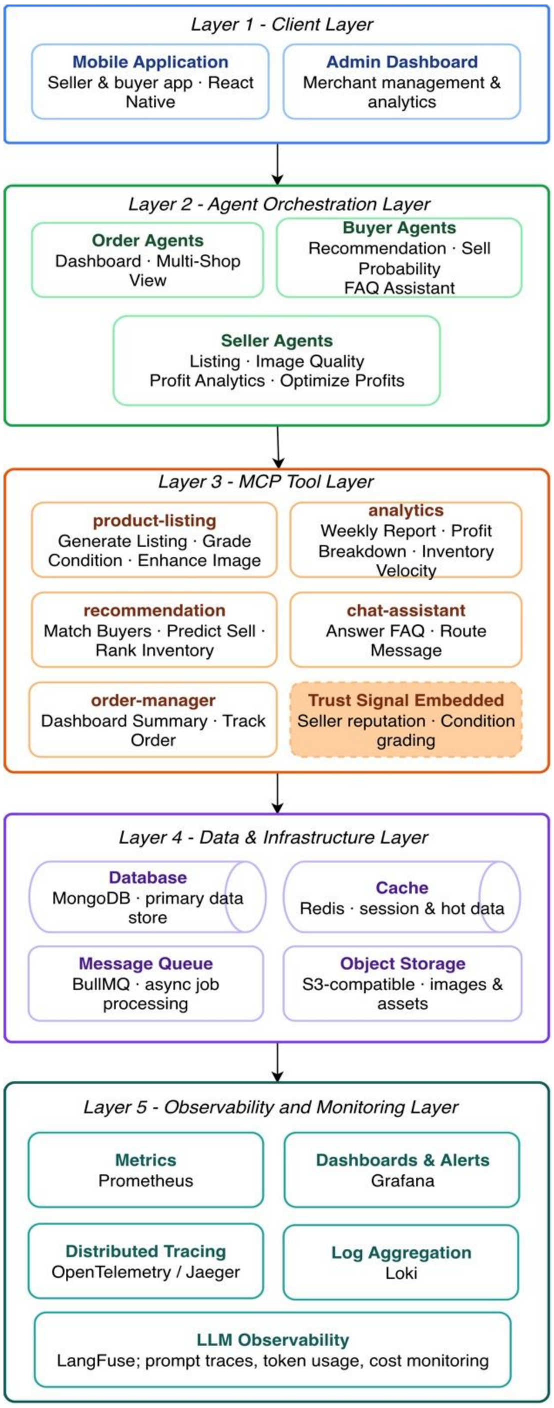

Figure 4. Five-Layer System Architecture

6.1 MCP Tool Servers

Five MCP servers wrap bounded domains of the backend API. Development priority follows descending survey mean scores: the analytics server is first priority, driven by the highest merchant demand for sales intelligence and profit visibility. The order-manager and recommendation servers follow, reflecting strong interest in multi-store management and buyer matching. The product-listing server carries the heaviest trust responsibility, as AI condition grading and product history are the two most actionable trust signals identified in the literature [4].

Table 2. MCP Tool Servers — Survey Score Range and Exposed Tools

MCP Server	Score Range	Tools Exposed
analytics	4.60 – 4.73	weeklyReport · profitBreakdown · inventoryVelocity · optimiseProfit
order-manager	4.20 – 4.67	dashboardSummary · trackOrder · multiShopView
recommendation	4.27 – 4.67	matchBuyers · predictSellProbability · rankInventory
chat-assistant	4.20 – 4.47	answerFAQ · routeMessage
product-listing	3.53 – 4.00	generateListing · enhanceImage · gradeCondition

6.2 MCP Interaction Flow

Three interaction flows illustrate how MCP tools compose within the agent layer [3], as shown in Figure 5, using server: tool notation, where the prefix names the MCP server and the suffix its exposed tool. On product search, the Recommendation Agent calls recommendation: matchBuyers against Qdrant embeddings [17] to obtain a ranked shortlist, then fans out recommendation: predictSellProbability in parallel — annotating each result with a sell-probability score so the highest-confidence matches surface first without a separate ranking pass.

Inbound buyer messages follow a two-step fallback: the Chat Assistant Agent queries chat-assistant: answerFAQ and, when confidence falls below threshold, escalates to chat-assistant: routeMessage, forwarding the conversation to the seller's unified inbox via the order-manager server. This isolates automated responses from low-confidence cases, preserving the direct negotiation dynamic merchants elected not to automate.

On sale confirmation, the Order Dashboard Agent calls order-manager: trackOrder and order-manager: multiShopView atomically in a single A2A round-trip [13], eliminating per-channel polling. The analytics server consumes the same event asynchronously via BullMQ, ensuring weeklyReport and profitBreakdown always reflect settled orders without additional seller input.

Across all three flows, agents call MCP tools directly rather than routing through a central orchestrator, keeping latency low and allowing each MCP server to evolve, scale, or fail independently without cascading changes through the rest of the system.

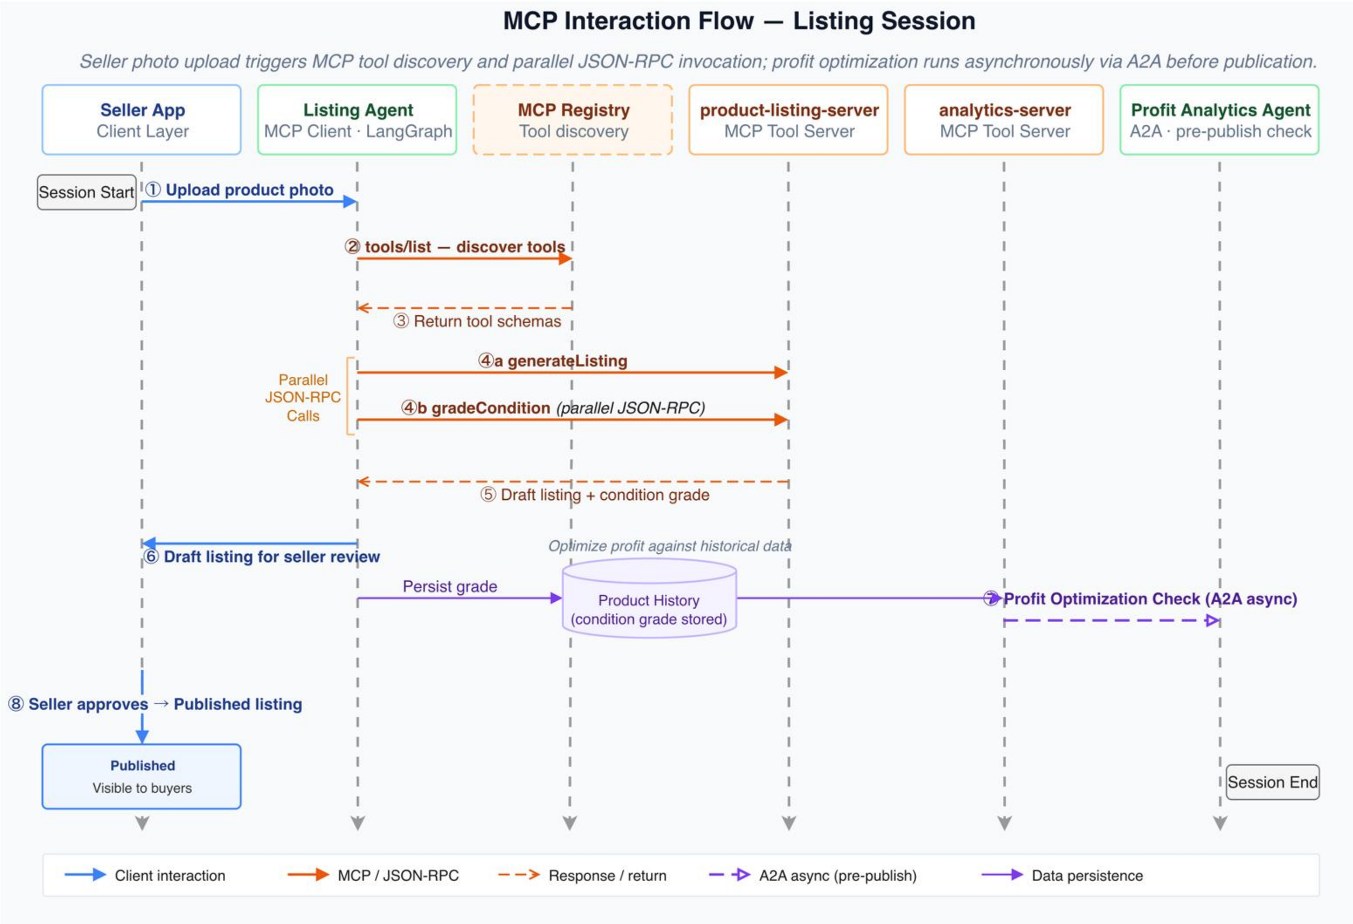

Figure 5. Listing Agent MCP Interaction Sequence

6.3 Trust Layer Design

Grounded in signaling theory findings [4], the trust layer addresses the three components that most strongly influence buyer trust in secondhand commerce (Table 3):

Table 3. Trust Layer Components

Component	Implementation	Signal Type (per [4])
AI Condition Grading	product-listing: gradeCondition	Product history / refurbishment info
Seller Reputation Score	Order completion rate · reviews · response time	Seller reputation (strongest predictor)
Verified Seller Badge	KYC check at onboarding	Institutional trust signal

Three trust components are architecturally specified, but each requires targeted empirical research before production implementation. First, the verified-seller badge and KYC workflow present a Myanmar-specific design challenge: limited formal identification infrastructure means many merchants operate without business registration, requiring a tiered verification approach — phone-verified, NRC card-verified, and business-registered — whose trust-building threshold has not been established for Myanmar buyers. Second, the seller reputation scoring algorithm synthesizes order completion rate, response time, and review scores into a composite signal, but no published study has calibrated which input metric most strongly predicts perceived seller credibility in Myanmar's social commerce context [2, 4]. Third, AI condition grading (product-listing: gradeCondition) requires buyer-side threshold calibration: the minimum model confidence at which an AI-generated grade raises purchase intent — rather than introducing doubt relative to the seller's own description — is unknown for this user population [20]. Fourth, the three signals interact: a high condition grade from a low-reputation seller may cause buyers to discount both, and this interaction effect has not been measured. Conjoint analysis across Myanmar buyer segments is needed before the trust layer weighting can be finalized.

6.4 Multi-Agent Design

Four agent families are implemented as LangGraph state machines coordinating via A2A (Table 4):

Table 4. Multi-Agent Families and Roles

Family	Agents
Seller	Listing · Image Quality · Profit Analytics
Buyer	Recommendation · Sell-Probability · FAQ Assistant
Order	Order Dashboard · Multi-Shop
Governance	Fraud Detection · Compliance (future scope)

6.5 Technology Rationale

Table 5 lists the primary architectural technologies and the design decision motivating each choice.

Table 5. Technology Rationale by Layer

Layer	Technology	Rationale
Agent orchestration	LangGraph [15] · LangChain [16]	Provides stateful, cyclical agent workflows with explicit state transitions, enabling reproducible and auditable agent behavior.
MCP protocol	Anthropic MCP SDK [3] · JSON-RPC 2.0	Standardizes tool discovery and invocation between AI agents and backend services; each MCP server can be versioned and deployed independently.
Agent messaging	Google A2A Protocol [13]	Provides a standard inter-agent communication layer, allowing Seller, Buyer, and Order agents to coordinate without direct coupling.
Vector retrieval	Qdrant [17] — product embeddings	Supports semantic search over product embeddings, enabling AI-assisted product discovery beyond keyword matching.
AI observability	LangFuse [18] — agent tracking	Traces prompt automated intelligent actions like chains and token usage per agent call, providing cost visibility and regression detection when models are updated.
MLOps	MLflow [19] — model registry	Manages model versions for sell-probability and pricing models, enabling rollback when production performance degrades.
Image processing	YOLO (ONNX Runtime, quantized) · Rembg — object detection + background removal	YOLO models exported to ONNX Runtime and quantized for low-precision inference give real-time object detection at low latency and low bandwidth, matching the entry-level mobile devices and constrained data plans typical of Myanmar's price-sensitive merchant base; Rembg removes product backgrounds for consistent, professional-looking listing images without requiring merchants to reshoot photos.
Async inference	BullMQ workers (extended for AI jobs)	Decouples long-running AI inference tasks from the HTTP request cycle, preventing timeout errors on slow model calls.

7. DISCUSSION

The survey-driven methodology ensures MCP server development effort is proportional to demonstrated merchant need — features with low survey interest are deferred rather than built speculatively. The evidence chain (survey mean score → server priority) provides an auditable justification that strengthens academic validity and stakeholder buy-in, embedding trust infrastructure as a first-class architecture requirement — rather than a post-launch enhancement — and responds directly to the empirical finding that trust is the dominant purchase driver in Myanmar's social commerce market [2]. Platforms that address operational efficiency without addressing trust risk create a fast but untrustworthy marketplace that fails to convert Myanmar's Facebook-native buyer base.

MCP's composability provides a practical advantage for resource-constrained teams: condition grading can be upgraded to a more accurate computer vision model [20] without touching any other agent or MCP server, allowing trust signal quality to improve iteratively post-launch. The A2A messaging layer further ensures that multi-agent coordination remains loosely coupled and independently auditable — a property that is particularly important for the trust layer's reputation scoring, where scoring logic changes must not inadvertently affect listing or pricing agents.

The trust findings from prior Myanmar social commerce research [2, 4, 9] are reflected in the architecture through three mechanisms: AI condition grading (product-listing-server: gradeCondition), seller reputation metrics exposed via analytics-server, and product provenance tracking. These are treated as requirements-driven components derived from survey evidence and supporting literature, not as optional additions.

8. LIMITATIONS AND FUTURE WORK

This proposal is informed by a pilot needs-assessment survey involving fifteen merchants. While qualitative richness justifies the design direction, generalisability of mean scores is limited pending a larger survey across multiple Myanmar cities, product categories, and seller experience levels. The trust layer components (seller reputation scoring, verified-seller badge, KYC) are currently defined at a high level; their detailed design requires empirical study specific to Myanmar's limited formal identification infrastructure.

Future directions fall into two groups. Trust layer research: (i) KYC feasibility study establishing which identity documents are realistically available to Myanmar merchants and what verification tier triggers sufficient buyer trust; (ii) seller reputation weighting validation — a buyer-side survey calibrating whether order completion rate, response time, or review score is the dominant credibility signal in Myanmar's social commerce context, extending Negash and Akhbar [4]; (iii) AI condition grade threshold calibration — a buyer experiment measuring the model confidence level at which an AI-generated condition grade improves purchase intent over a seller self-description, extending He et al. [20] to a trust context; and (iv) multi-signal conjoint analysis measuring how Myanmar buyers weigh combinations of condition grade, reputation score, and verified-seller badge against each other, including interaction effects for mismatched signals. Architecture research: (v) A/B evaluation of MCP-generated versus manually written listings on buyer conversion rate; (vi) knowledge-graph RAG extension (Neo4j [21]) for semantic product relationships; (vii) federated learning for privacy-preserving demand forecasting across merchant cohorts; (viii) architecture adaptation for cross-border Myanmar–Thailand seller corridors; and (ix) Burmese-language LLM benchmarking for listing quality and FAQ accuracy.

9. CONCLUSION

This paper makes four primary contributions. First, it identifies the operational requirements of Myanmar secondhand and local fashion merchants through a structured needs-assessment survey designed to capture pain-point severity and feature priorities. Second, it maps empirical requirements directly to an MCP-driven multi-agent enterprise architecture, with development priority determined by survey evidence. Third, it integrates trust mechanisms — condition grading, seller reputation, and product provenance — as first-class architectural components, grounded in prior findings from Myanmar social commerce research. Fourth, it provides a practical architectural blueprint for AI-enabled digital transformation in

Myanmar's social-commerce ecosystem, positioned for deployment on an existing production platform. Beyond this initial scope, future work extends the architecture along two fronts (Section 8): empirical validation of the trust layer — calibrating seller-reputation weighting, verified-seller KYC thresholds, and AI condition-grade confidence against actual Myanmar buyer behavior — and platform growth through knowledge-graph-based product retrieval, federated demand forecasting across merchant cohorts, cross-border Myanmar–Thailand seller support, and Burmese-language LLM benchmarking.

10. ACKNOWLEDGMENTS

The authors thank the Myanmar merchant community members who participated in the survey design review, and the anonymous reviewers whose feedback strengthened the methodology, trust architecture, and Myanmar context sections.

11. REFERENCES

[1] DataReportal. 2024. Digital 2024: Myanmar — Global Digital Insights. Retrieved from https://datareportal.com/reports/digital-2024-myanmar

[2] Aung, T., Liana, S. R., Htet, A., and Bhaumik, A. 2024. Analyzing the determinants of consumer buying behavior on Facebook Marketplace in Myanmar: a quantitative approach. J. Electr. Syst. 20, 9s (2024), 1142–1152. Retrieved from https://journal.esrgroups.org/jes/article/view/4473

[3] Anthropic. 2024. Model Context Protocol Specification. Retrieved from https://modelcontextprotocol.io

[4] Negash, Y. T. and Akhbar, T. 2024. Building consumer trust in secondhand fashion: a signaling theory perspective on how consumer orientation and environmental awareness shape engagement. Clean. Responsible Consum. 14 (2024), 100211. DOI:https://doi.org/10.1016/j.clrc.2024.100211

[5] Ricci, F., Rokach, L., and Shapira, B. (Eds.) 2015. Recommender Systems Handbook. 2nd ed. Springer, New York.

[6] Bae, Y., Choi, J., Gantumur, M., and Kim, N. 2022. Technology-based strategies for online secondhand platforms promoting sustainable retailing. Sustainability 14, 6 (2022), 3259. DOI:https://doi.org/10.3390/su14063259

[7] Economic Research Institute for ASEAN and East Asia (ERIA). 2021. Digital Economy of Myanmar. In Digital Economy Integration in ASEAN, chap. 6. ERIA, Jakarta, Indonesia. Retrieved from https://www.eria.org/uploads/6_ch_6-Myanmar.pdf

[8] Kyi, N. P. W. 2024. An empirical study of information search and alternative evaluation, purchase and post purchase affecting customer satisfaction and e-commerce business development in Myanmar. M.S. thesis, Bangkok Univ., Bangkok, Thailand.

[9] Lita, R. P., Meuthia, M., Rahmi, D. Y., and Syafrida, M. F. 2024. Does social presence on social commerce platform attract buying intention of Indonesian local food? J. Appl. Eng. Technol. Sci. 6, 1 (2024), 174–191. DOI:https://doi.org/10.37385/jaets.v6i1.5852

[10] Krishnan, N. K. 2025. Advancing multi-agent systems through Model Context Protocol: architecture, implementation, and applications. Preprint arXiv:2504.21030. Retrieved from https://arxiv.org/abs/2504.21030

[11] Chen, K., Sun, Y., Keung, J., Mao, Z., and Ma, X. 2026. Understanding how enterprises adopt the Model Context Protocol for LLM-driven software engineering. Preprint arXiv:2606.09182. Retrieved from https://arxiv.org/abs/2606.09182

[12] Wooldridge, M. 2009. An Introduction to MultiAgent Systems. 2nd ed. Wiley, Chichester, UK.

[13] Google. 2024. Agent-to-Agent (A2A) Protocol Documentation. Retrieved from https://google.github.io/A2A/

[14] Witkin, B. R. and Altschuld, J. W. 1995. Planning and Conducting Needs Assessments: A Practical Guide. SAGE Publications, Thousand Oaks, CA.

[15] LangChain Inc. 2024. LangGraph Documentation. Retrieved from https://langchain-ai.github.io/langgraph/

[16] LangChain Inc. 2024. LangChain Documentation. Retrieved from https://python.langchain.com/

[17] Qdrant Solutions GmbH. 2024. Qdrant Documentation. Retrieved from https://qdrant.tech/documentation/

[18] Langfuse GmbH. 2024. LangFuse Documentation. Retrieved from https://langfuse.com/docs

[19] Databricks Inc. 2024. MLflow Documentation. Retrieved from https://mlflow.org/docs/

[20] He, K., Zhang, X., Ren, S., and Sun, J. 2016. Deep residual learning for image recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). 770–778.

[21] Neo4j Inc. 2024. Neo4j Documentation. Retrieved from https://neo4j.com/docs/