dmi:

# ADMC 2026

Academic Design Management Conference

# The 25th dmi: Academic Design Management Conference Proceedings

Hosted At

Georgia Tech

11-12 AUGUST 2026
GEORGIA INSTITUTE OF TECHNOLOGY
ATLANTA, GA. USA
WWW.DMI.ORG/ADMC

dmi: ADMC

## Conference Organization

**Carole Bilson**, *President, DMI*
**Chris Hancock**, *ADMC 2026 Conference Program Coordinator*

## DMI Board of Directors

**Ernesto Quinteros (Chair, DMI Board),** *Chief Experience Officer, Korrus*
**Carole Bilson,** *President, DMI*
**Kelly Jura (Chair, DMI Advisory Council),** *VP, Brand and UX, ScreenPal*
**Jerry Kathman,** *dmi: Life Fellow, Emeritus Chair DMI Board, Founder & Chairman, LPK*
**Betsy Goodrich, (Secretary),** *VP/Co-Founder, MANTA Product Development Inc.*
**Jill Lawrence,** *VP Design, Crown Equipment Co*
**Joe Perdew,** *VP Store Design, Target*
**Shani Sandy,** *VP, Experience Design, Tech Lifecycle Svcs, IBM*

## Hosted At

25th DMI: Academic Design Management Conference
Leading Design & Innovation Research & Education in the Age of AI
Atlanta, GA, USA, 11-12 August 2026

# A Service-AI Framework for Human-Centered Design in Enterprise Sales: Orchestration, Execution, and Roll Out

Parag DESHPANDE [a], Yingjie LI [a]
[a] Docusign

*This paper investigates the Design with AI paradigm in Docusign's enterprise sales, focusing on mitigating the 40-70% administrative burden faced by Account Executives and Customer Success Managers. The primary objective is to establish a Service-AI framework in which retrieval-augmented orchestration engines handle high-velocity data synthesis, freeing humans to focus on strategic decision-making. Legacy enterprise quoting tools are often rigid, forcing sales professionals to use disconnected spreadsheets. While current research focuses on speed, a theoretical gap exists in understanding how contextual intelligence supports non-standard, creative deal structures. We ask: How can AI-driven contextual intelligence improve user agency and operational efficiency in high-stakes financial decisions without creating a black-box dependency?*

*Using a six-phase Experience Architecture methodology, this research mapped the Quote-to-Cash lifecycle, revealing manual workarounds used to bypass legacy systems. A rapid pretotyping and Minimum Lovable Product (MLP) pilot with 37 Docusign participants yielded a 9.5/10 user rating, raised baseline usability scores from 1.6 to 3.1 (out of 4), and achieved a 72% automated proposal-generation rate. Operationally, the system projects an estimated 85,000+ annual hours saved, equivalent to 82 Full-Time Equivalents (FTEs). The research concludes that Design with AI is most effective as a symbiotic loop, balancing machine logic and scale with essential human empathy and accountability.*

***Keywords****: Design with AI; Service-AI Framework; Human-AI Synergy; Experience Architecture; Cognitive Augmentation; Enterprise Design; User Agency*

## Introduction

### *The Enterprise Crisis*

In contemporary B2B and enterprise ecosystems, administrative overhead increasingly overburdens sales professionals. Our qualitative research and current market data indicate that Account Executives (AEs) tasked with new business acquisition and Customer Success Account Managers (CSAMs) dedicated to client retention and renewals spend between 40% and 70% of their time on non-revenue-generating administrative tasks. As one AE bluntly summarized the friction with legacy platforms, "We spend more time fighting the tech than actually selling." Poor data hygiene, manual data transfers across disparate systems, and the fragmented nature of modern tech stacks further exacerbate the issue. Users open multiple tabs and toggle between CRM systems, email, analytics dashboards, and spreadsheets. Thus, it creates cognitive overload, leading to error-prone estimates and lost revenue opportunities.

### *Problem Statement*

Legacy enterprise quoting tools (defined here as rigid, heavily siloed early-generation SaaS or on-premise architectures) used by AEs and CSMs are frequently described as rigid, outdated, and cumbersome. These monolithic architectures lack contextual intelligence and suffer from slow loading times, forcing sales professionals to rely on disconnected spreadsheet workarounds to maintain agility when constructing creative deal structures. This behavior exemplifies what Parasuraman and Riley (1997) identify as the "disuse" of automation, where users abandon rigid or opaque systems in favor of manual, higher-effort alternatives that offer greater cognitive transparency and control. Users describe standard enterprise reporting as convoluted, with multiple fields sharing the same name but having different meanings. Consequently, AEs spend up to an hour to build a single quote, relying heavily on Revenue Operations for manual intervention.

### *The Proposition*

This paper introduces the "Service-AI" framework as a strategic bridge between high-velocity data and nuanced human decision-making. We aim to transform the enterprise quoting experience by applying the principles of Composable AI. This architectural approach allows independent, modular artificial intelligence capabilities (e.g., retrieval, generation, and logic engines) to be combined like building blocks to execute complex workflows (Workato, 2026). This is particularly critical in the modern "everything-as-a-service" (XaaS) economy, where Configure, Price, Quote (CPQ) tools have evolved from back-office utilities into the primary engines of strategic growth. We argue that the enterprise must transition away from brittle, deterministic decision trees toward AI-powered "reasoning engines." By moving away from a monolithic "black box," this framework shifts quoting from manual data entry to an AI-assisted process, saving about 15 to 20 minutes per opportunity. In this context, we define a black box simply as UX opacity: a situation where the user cannot trace how the system arrived at a specific business conclusion.

### *Research Questions*

A central inquiry drives our study: How can enterprise systems leverage AI to restore human agency without creating a dangerous "black-box" dependency? Furthermore, how can we design interfaces that support non-standard, creative deal structures while remaining transparent, accountable, and highly performant?

## Literature Review & Theoretical Grounding

### *The Enterprise UI Gap and The Paradigm Shift*

The rapid evolution of enterprise capabilities often outpaces organizational adoption. This friction is encapsulated by Martec's Law, which observes the widening gap between exponential technological growth and logarithmic organizational adaptation due to rigid legacy systems. Historically, the dominant paradigm in software development was Design for AI, where users acted merely as passive recipients of automated outputs. Yet, as Cico et al. (2025) demonstrate in their empirical study of generative AI workflows, applying advanced models to legacy architectures without transparent interaction mechanisms consistently triggers significant user resistance.

Bridging this gap requires modern composable enterprises to shift toward Design with AI, treating the algorithm as a collaborative partner rather than an isolated authority. This approach aligns with

Shneiderman's (2020) Human-Centered AI (HCAI) framework and is operationalized by the HAID (Human + AI in Design) framework (Mohammed, 2025). HAID posits that in enterprise settings, AI should absorb deterministic, high-volume tasks, thereby elevating the human worker from a manual executor to a strategic orchestrator.

### *The Technological Shift: Composable Architectures and Agentic RAG*

While contemporary research often prioritizes computational speed in sales automation (Smith, 2024; Lee, 2025), a theoretical gap remains in understanding how contextual intelligence supports non-standard, highly creative deal structures. Modern systems address this limitation through Composable AI and Agentic AI, which represent autonomous architectures that execute defined goals by dynamically routing tasks to specialized models.

For enterprise Account Executives, effective tooling must account for buyer sentiment and complex organizational hierarchies. Modular RAG (Retrieval-Augmented Generation) architectures decompose the retrieval process into specialized stages: Query Planners break down complex queries, Context Processors rank historical deals by relevance, and Validators ensure compliance. Building on this, Agentic RAG systems utilize the ReAct (Reason + Act) framework (Yao et al., 2022) to proactively fetch missing information. Transitioning from static databases to these intelligent reasoning engines provides the technical foundation necessary to handle quoting variables at scale.

### *Buyer Psychology and The Transparency Imperative*

Despite these technological capabilities, AI adoption in high-stakes financial decisions is heavily mediated by user psychology. In B2B pricing, cognitive biases play a crucial role. Interfaces must effectively leverage the anchoring effect through guided selling or the decoy effect via tiered pricing models to provide buyers with a necessary sense of control. When users cannot trace how an AI arrived at a recommended price, they often experience "black-box alienation," which drives algorithmic aversion (Ribeiro et al., 2016).

To mitigate this opacity, Panchal (2025) identifies "Shared Autonomy Controls" and "Explainability on Demand" as mandatory design patterns for next-generation Agentic UX. In our work, we draw on the MATCH framework (Modular Architecture for Transparent and Controllable Human-AI interaction) (Vanbrabant et al., 2025) to map underlying logic to structural building blocks. By exposing these logic blocks through interactive UI components, which are conceptually similar to explanatory models like SHAP (Lundberg & Lee, 2017) or LIME, the system provides post-hoc visualizations of model behavior. This design strategy directly supports Hutchins' (1995) theory of distributed cognition, in which the machine manages multi-system data retrieval while the human manages relationship context and ethical accountability.

### *Synthesis: Towards a Service-AI Framework*

Ultimately, while Composable AI and Agentic RAG offer unprecedented processing power, their successful deployment in enterprise sales depends on overcoming the psychological barriers of black-box opacity. Current literature indicates a pressing need for an architecture that combines the operational velocity of autonomous agents with the transparent, human-in-the-loop oversight required for high-stakes deal-making. This specific synthesis forms the theoretical basis for our proposed Service-AI Framework.

## Methodology: The Six-Phase Experience Architecture

To bridge business strategy with user-centric design, this study employed a rigorous six-phase Experience Architecture approach, guided by a Bringing Order to Chaos philosophy.

### *Phases 1-2: Discovery & Journey Mapping*

We conducted persona-led contextual inquiry workshops with 18 participants across the Quote-to-Cash lifecycle, representing Account Executives (Install and NewCo), CSAMs, Deal Desk, and Revenue Operations. Gathering over 550 observational notes, we identified critical friction points: participants spent up to 70% of their time on manual estimates in highly complex, disconnected Google Sheets before migrating data into official CPQ systems. We also compared these findings with industry-wide benchmarks and matched them to the internal time-study data. This discovery phase was heavily influenced by the need to avoid the technology trap, a phenomenon responsible for a 70 to 88% failure rate in digital transformations, and to focus instead on solving real users' real pain points (Cornelissen, 2025). By adhering to a Humans × Machines convergence model, we intentionally redesigned roles and workflows to integrate human judgment with machine speed, rather than simply superimposing modern technology onto idiosyncratic and behaviorally entrenched processes. The Journey Map (JM) exercise proved invaluable, bringing cross-functional teams onto a single page to view the process from the user's perspective. It clarified the overlaps and hand-offs between various personas. A significant benefit of the JM was the discovery of surprising user sentiments and unexpected "hot spots" for pain points that we were then able to prioritize.

### *Phase 3: The Experience Architecture*

Along with the qualitative research, we analyzed complex data flows from 350,000 historical opportunities to pinpoint areas for maximum cognitive relief. This phase defined the architectural requirements for a system capable of dynamically fetching context, verifying permissions, and providing logical recommendations without disrupting existing sales workflows. We established a modular experience architecture that can be scaled up or down as needed, serving as the blueprint for all business and technology discussions. This architecture incorporates Composable AI components and flows designed to be highly adaptable for future requirements.

### *Phases 4-5: Pretotyping & Qualitative Validation*

To rapidly validate user desirability before writing production code, we employed Figma-based pretotyping (Alberto Savoia, 2019) with 9 participants to assess receptiveness to algorithmic recommendations, one-click proposal generation, and human-control override features. Crucially, this phase focused on human control, enabling users to override system-generated artifacts. As we continued developing the application, we iteratively tested it to collect feedback from real users. Eventually, we conducted 1-on-1 qualitative usability testing in a staged sandbox environment with 7 users. The baseline and current quoting experience were rated a dismal 1.6 out of 4, primarily due to slow load times, lack of data granularity, and spiderweb reporting. User testing demonstrated increased satisfaction, with an average rating of 3.1 out of 4, and participants highlighted the consolidation of information as a major win. Thematic analysis revealed four core user needs: Efficiency & Automation, Strategic Enablers, Customization, and Transparency & Control.

### *Phase 6: Pilot & Global Rollout*

We released an MLP (Minimum Lovable Product) as a pilot to 37 participants in the US and LATAM for real-world telemetry. The pilot phase allowed us to continue collecting feedback and to ensure that the crucial backend services continued to work at the required scale and performance. After one month of the pilot phase, adoption was promising, with 21 active users creating 438 estimates and converting them into 109 quotes. During this phase, they also generated 52 slide decks. In parallel, we worked with the enablement team to strategize the general availability rollout to over 700 users in North America and, later, to 1,800 users globally.

## The Service-AI Framework

The Service-AI Framework is a modular architecture that integrates AI and human labor to solve the "Scalability-Personalization Paradox," enabling high-volume service without losing nuanced human touch (Figure 1). By standardizing communication and delegation protocols, the framework creates a highly replicable system that can be ported across diverse industries to handle complex workflows. It is theoretically "symbiotic" because it establishes a mutualistic relationship in which the retrieval and generation layers handle the computational "heavy lifting," while humans provide the essential ethical and emotional "nervous system." This co-evolutionary partnership ensures that the unit’s combined output far exceeds the capabilities of either entity working in isolation. Three interconnected pillars define it:

### *Orchestration*

The orchestration layer acts as a modular synthesis engine, pulling real-time consumption data, historical contracts, and firmographic intelligence from disparate systems. The Model Context Protocol

(MCP) supports the Service-AI framework to serve as a standard "USB-C for AI," enabling the agentic environment to interact with external tools and databases without custom security or ad hoc integrations.

### *Execution*

The Quick Quote interface allows for rapid iteration. To address latency issues inherent in distributed multi-agent systems, the architecture employs stack-agnostic optimizations, such as prompt caching (reducing latency by up to 80% by storing repeated instructions) (AI Lab, 2025), to ensure real-time UI responsiveness during complex deal modeling.

### *Human Override*

Human intervention is the critical "Agency" layer, treating the generative model as a "Strategic Co-pilot" rather than as an autonomous underwriter. The framework ensures that AI makes recommendations and the human user retains ultimate control to validate, discount, or edit model-generated artifacts.

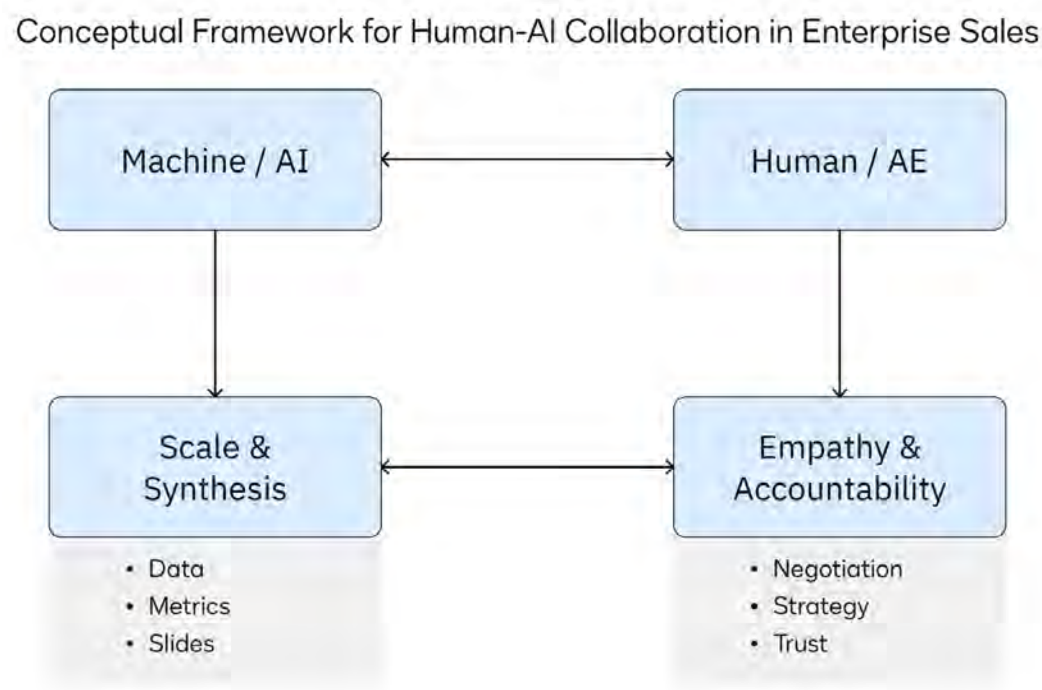


*Figure 1 The Service-AI Symbiotic Loop*

The Service-AI framework is fundamentally grounded in several core theoretical precepts. Parasuraman and Riley (1997) identified two primary modes of human-automation failure: excessive reliance, termed misuse, and insufficient reliance, termed disuse. These failures commonly stem from a lack of clarity or predictability in the system's behavior. The Service-AI framework directly mitigates these risks by incorporating a "human override" function and an "approval gate" structure. Furthermore, Parasuraman, Sheridan, and Wickens (2000) categorized human tasks amenable to automation into four distinct functions: information acquisition, analysis, decision selection, and action implementation. Their work emphasizes the critical importance of judiciously selecting the level of automation for each function to ensure effective human-system collaboration.

The design of Quick Quote strongly aligns with the concept of human-AI collaboration: the Retrieval-Augmented Generation (RAG) pipeline independently executes information gathering and review, suggesting next steps, but requires human initiation for the final action. This approach incorporates several of the 18 common-sense design rules for human-AI interaction proposed by Amershi and colleagues (2019), such as ensuring the relevance of suggestions, allowing corrections, and clearly indicating the predictive model's confidence levels. However, the Service-AI framework expands on these ideas by addressing a more comprehensive set of challenges related to responsible, company-wide human-AI teamwork. Specifically, it addresses the "bigger picture" management issues, including system decomposition, transparency, required approval workflows, and orchestrating multiple AI components, aspects that other frameworks do not fully resolve.

## Implementation: Quick Quote Case Study

### *User Persona Focus*

We centered our design intervention on the AE and CSM personas (Figure 2). For the AE, the main goal was to streamline quote creation and deal velocity; for the CSM, the goal was to shift from reactive support to proactive renewal orchestration. The user persona studies helped us understand the roles and responsibilities, as well as their overlaps. The personas were imperative as we explored concepts and designs to help improve collaboration.


*Figure 2 User Personas*

### *Unifying Data & Usage Trends*

Previously, users had to check information across 10 tabs in different browsers and copy it manually. Quick Quote consolidates this into a single pane and automatically calculates discounts and net prices upon data entry (Figure 3). A highly celebrated feature was the Usage Trends & Details module, which integrates consumption data, feature adoption, and "sticky features" into a clear, unified view, enabling reps to easily justify renewals and upsells.

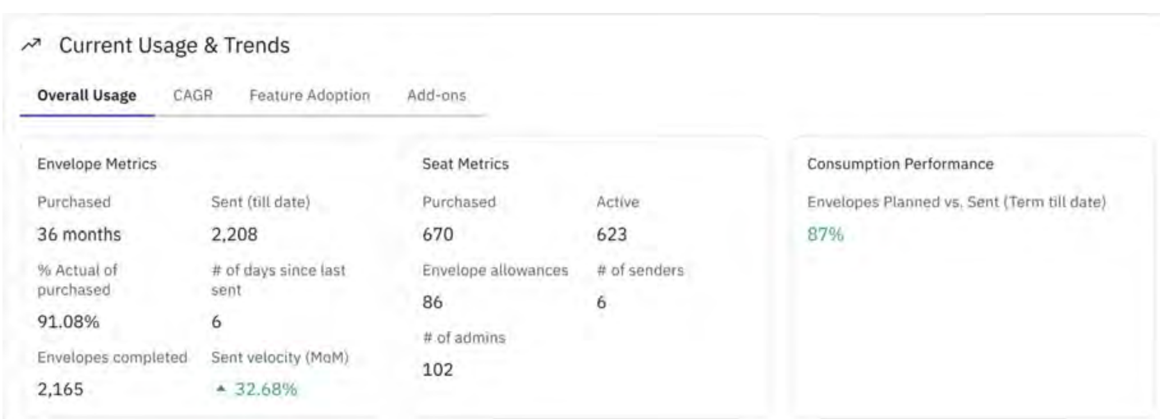


*Figure 3 Snapshot of unified data and usage trends from the Quick Quote application*

### *Feature Highlight: The Proposal Deck & Scenario Modeling*

One of the most profound design interventions was the automated generation of the proposal deck. Historically, users manually collated infographic data and pricing in Google Slides, a process prone to human errors. Our modular AI architecture enabled "one-click" generation of customer-ready presentations by collating data from multiple sources and unifying it to enable great storytelling. Furthermore, the "Clone & Modify" feature provided an essential audit trail, allowing reps to spin up side-by-side comparative estimates for sophisticated scenario modeling instantly and automatically highlight price deltas. (Figure 4)

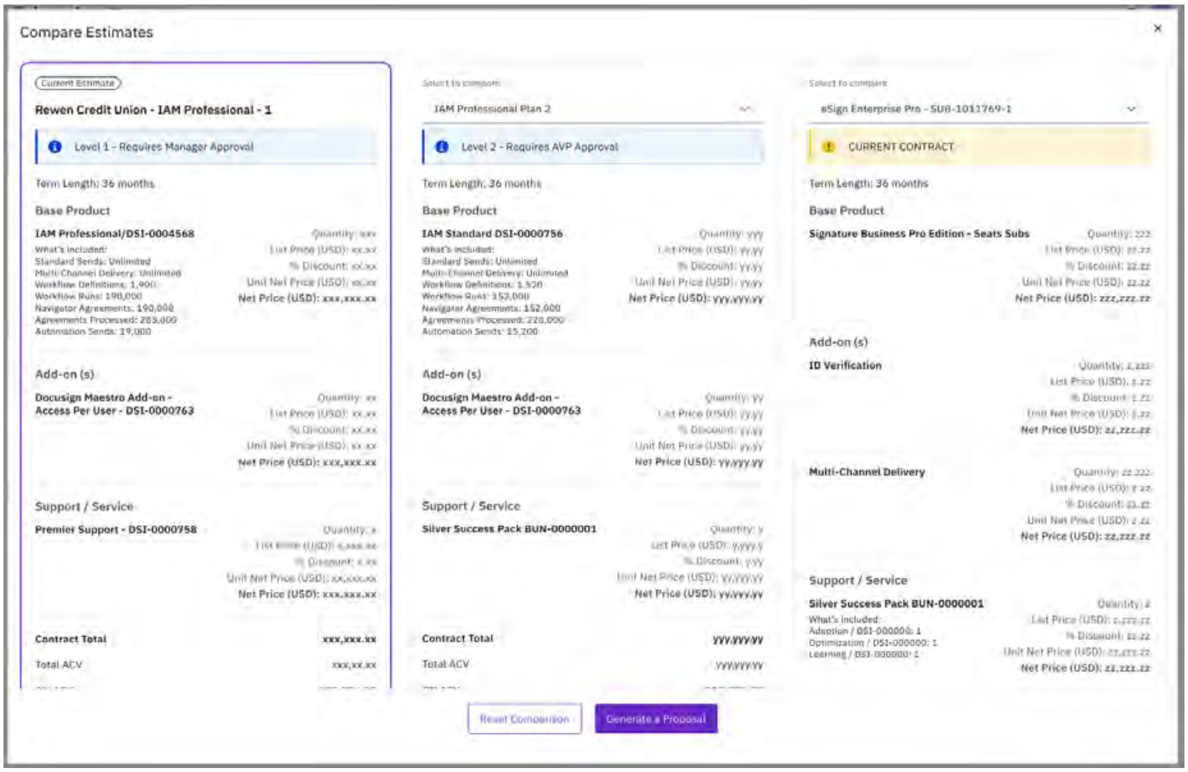


*Figure 4 Snapshot of compare and scenario modelling from the Quick Quote application*

### *Design Intervention: From Manual Hack to AI-Assisted Strategy*

By auto-populating admin contacts and payment terms based on previous contracts, Quick Quote eliminated frustratingly repetitive tasks. Instead of performing manual data entry across disconnected spreadsheets, the user now reviews pre-filled configurations, adjusts global discount variables, understands approval matrices in real time, and approves the strategy. The dynamically generated proposal deck is fully editable, giving the user 100% control over adding or modifying content before easily and comfortably sharing it with customers.

### *Designing for AI Transparency and Trust*

This commitment to user-driven iteration was equally evident in our granular design decisions, particularly regarding the transparency of the underlying algorithms. During the configuration stage, the system is designed to prevent users from starting from a blank slate by pre-filling recommended product quantities based on actual customer usage data. Alongside this, the system provides smart pricing discounts. Initially, the interface displayed these machine-learning (ML) driven pricing suggestions using a standard information icon that showed brief explanatory text when hovered over. However, observational research revealed that users hesitated to trust the system; they explicitly wanted to understand exactly how the recommended discount was calculated before presenting it to a client. In response, the design team replaced the generic tooltip with a prominent, dedicated AI visual component. When users interact with the new AI icon, it reveals a comprehensive explanation card. (Figure 5) By increasing the typographic scale of the discount value, employing a subtle gradient background to visually signify an intelligent insight, and providing an explicit explanation stating that the recommendation is based on an analysis of historical data and similar customer deal benchmarks, we transformed user skepticism into immediate strategic confidence.

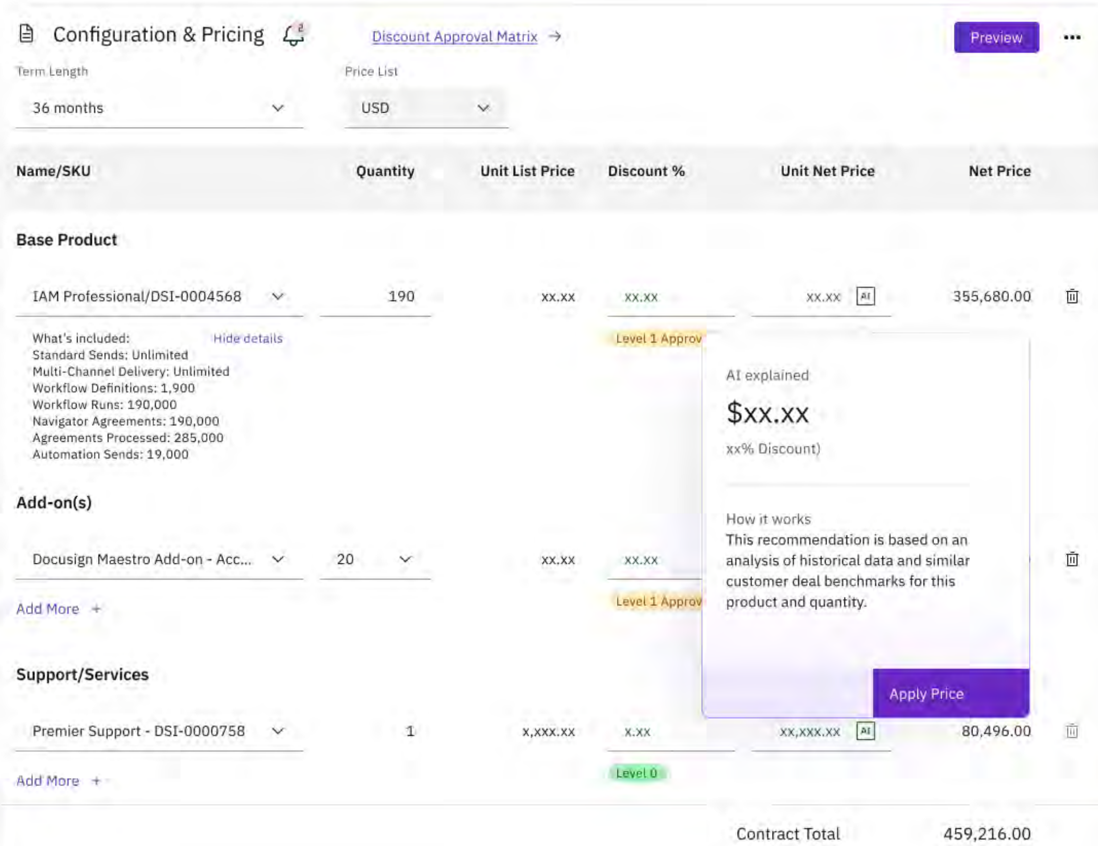


*Figure 5 Design with AI for trust and transparency, illustrative snapshot of AI pricing recommendation from the Quick Quote application*

# Results & Analysis

## *Quantitative Success*

The empirical data gathered during the pilot phase definitively validated the Service-AI framework:

- User Rating: The MLP achieved an exceptional **9.5/10 average user satisfaction** rating, with baseline user testing (usability) jumping from **1.6 to 3.1** (out of 4).
- Conversion Metrics: During the 1-month pilot phase, users loaded **111 unique opportunities** into Quick Quote, creating **438 estimates** and converting **109 of those into quotes**. During this phase, users also generated **52 automated slide decks**.
- Business ROI: Based on time-motion analysis and historical sales-opportunity data, the Service-AI framework delivers **85,000+ annual hours** in operational savings. These time savings restore two full weeks of productivity (54 weeks of output in a 52-week year) per user annually. By digitizing this effort, the organization adds the equivalent of **82 Full-Time Equivalents** (FTEs) without hiring a single new resource, projecting an **additional $10–$15 million in revenue** using the same team strength.

### *Qualitative Feedback*

The qualitative shift in user morale was equally significant. Participants consistently reported a considerable reduction in perceived cognitive load. Participants frequently expressed relief from cognitive friction, with one user noting the interface was 'pretty life changing.' At the same time, another highlighted the stark contrast with legacy systems: “If I never have to see it [the previous quoting tool] again... I'm sold.” A Strategic AE summarized the interface shift: "I just wanted to say it is a really great product... it can be very, very helpful to our day-to-day tasks, so I am really excited about it." The insistence on human-override mechanisms, such as the ability to customize AI-generated presentations, successfully mitigated "black-box" anxieties, thereby engendering deep user trust.

### *Post-Pilot Global Scale and Real-World Telemetry*

Between the initial submission of this manuscript and the final revision, the framework transitioned from a limited pilot phase into a full production environment. The system was deployed globally for direct deals to 1,750 active users across North America, EMEA, APJ, and LATAM. This rapid expansion provided a robust dataset to validate the system at an enterprise scale and monitor the transition from legacy architectures. Current telemetry demonstrates substantial operational adoption. At present, 67 percent of all human-generated quotes are eligible for automated processing through the Quick Quote architecture. Furthermore, strategic feature releases scheduled for the summer aim to increase this eligibility threshold to a 90 percent target.

Velocity and adoption metrics from recent production cycles validate the system's efficiency in high-stakes environments. During the most recent week of monitoring, the platform successfully processed 53 percent of all eligible enterprise quotes. The performance is particularly notable within the critical domain of business preservation, where the architecture accommodated over 71 percent of eligible on-time renewals.

The extensive telemetry also highlights nuanced adoption patterns across different business segments and geographic regions. For instance, the Commercial Business Unit demonstrates consistently higher adoption rates compared to the Enterprise Business Unit. This variance aligns with our expectations, given that commercial segments typically handle higher transaction volumes and have lower structural complexity. Similarly, regional data indicate that while adoption has steadily increased alongside global rollouts, specific markets such as APJ exhibit unique systemic friction points.

Critically, these quantitative insights directly inform our continuous design iteration process. To address lower adoption rates in complex enterprise and early-renewal scenarios, upcoming software updates will introduce advanced capabilities, including dedicated estimate-comparison interfaces, backdating features, and sophisticated credit calculators. Likewise, localized adjustments to pricing uplift handling are actively being developed to better support the APJ user base. This ongoing synthesis of global telemetry and targeted UX enhancement reinforces the core premise of our Service-AI framework: enterprise tools must continuously adapt to support intricate human workflows.

## Discussion: Human-AI Synergy

### *Restoring Agency*

The true success of the Service-AI framework lies in its ability to restore user agency. Legacy tools reduced highly trained sales professionals to manual data-entry clerks. By abstracting the complexity of data integration, the orchestration and automation layers absorb the administrative burden. The framework elevates the user from a "manual executor" to a "strategic orchestrator," enabling them to focus entirely on human-to-human relationship-building and complex negotiations.

### *Maturing the Organizational Design Culture*

Beyond the immediate usability improvements and operational time savings, the development of Quick Quote represents a fundamental maturation in organizational design management. Historically, internal tool development within the enterprise followed a rushed, ad-hoc, engineering-first trajectory. Product managers often base feature roadmaps on minimal foundational research, sometimes interviewing only two or three users before relying heavily on internal assumptions, prioritizing immediate product delivery over actual user utility. The Quick Quote initiative deliberately dismantled this approach by instituting a rigorous, continuous research cadence spanning the entire design lifecycle. By conducting more than five distinct rounds of user research, engaging more than 52 participants, and logging over 2,000 minutes of direct interview time, the design team established a culture of profound user advocacy. This extensive qualitative foundation enabled us to dynamically adjust product requirements and incorporate new functionalities as actual needs emerged in the field.

This responsive philosophy directly shaped our deployment strategy. Our initial roadmap outlined a Minimum Viable Product in October for a small cohort of 10 users, followed by a second release in November for 50 users, culminating in an exponential global-scale-up in March 2026 to over 1,800 users. While the team successfully met these overarching delivery milestones, the actual feature prioritization within each release remained highly fluid. For instance, when early participants highlighted a critical need for advanced usage tracking and Compound Annual Growth Rate metrics, we took that feedback directly back to the drawing board and integrated those precise capabilities into subsequent iterations. We made it a central tenet to listen to the sales field at every step and build what they genuinely required to succeed, rather than blindly executing a rigid, pre-determined schedule.

### *Organizational Impact and LLM Observability*

Scaling this architecture to 1,800+ global users implies a substantial shift in organizational capability. To maintain stability in a composable architecture, organizations must move beyond traditional uptime tracking and rely on robust LLM observability platforms. This ensures that teams can detect "prompt drift" and create feedback loops in which human corrections serve as learning signals for future model refinement. Implementing robust LLM observability platforms, such as LangSmith for trace visibility and Maxim AI for agent simulation, is critical for monitoring agent behavior in production settings.

### *Design Ethics, Standards, and Accountability*

In high-stakes enterprise sales, hallucinations or algorithmic bias carry high financial risks. The Service-AI framework operationalizes trust by enforcing "approval gates," human-in-the-loop checkpoints, and dynamic approval matrices that update in real-time based on discount percentages. Agentic models are programmed to admit uncertainty, codifying "deferral as a feature" to ensure that non-standard deviations are routed immediately to human Deal Desk specialists. Furthermore, aligning with global standards, specifically ISO/IEC 42001 for AI Management Systems, ensures that the development is ethically grounded, transparent, and legally accountable (Diab et al., 2025).

## Conclusion & Future Work

### *Summary*

The implementation of the "Quick Quote" MLP provides robust validation of the Service-AI framework. By abandoning rigid monolithic software in favor of composable, AI-driven modularity, organizations can eradicate systemic inefficiencies and embrace "Services-as-Software", buying measurable outcomes rather than static applications. The research proves that "Design with AI" functions best as a collaborative loop, combining machine scalability with essential human context.

### *Limitations*

While the empirical evaluation has expanded to a global deployment of 1,750 users, certain constraints remain. The current operational dataset intentionally excludes specific transaction types, notably auto-renewal deals, self-service direct deals, and partner-led motions. Furthermore, while adoption in regions like North America and EMEA is maturing rapidly, utilization in the APJ region remains comparatively low. This regional variance indicates that further localized system adaptations are necessary to fully accommodate unique regional pricing structures and uplift handling rules. Finally, while the initial rollout data strongly validates the Service-AI framework, observing the long-term impact of this paradigm shift on sales culture and employee retention will require multi-year longitudinal studies.

### *Future Research*

The design process does not end with a successful launch. As users recently identified the need for collaborative approval workflows, our team has already conceptualized the corresponding interfaces and scheduled them for validation in the coming weeks. Furthermore, future research must explore applying the Service-AI framework to other high-stakes, document-heavy enterprise domains, such as legal contracting and procurement. The previously rigid SaaS applications are now becoming composable multi-agent experiences. An open question for future research is whether sales motions can achieve true autonomy, operating without user initiation and requiring human-in-the-loop intervention only for critical decisions. In terms of technology, exploring low-level integrations like "Neurosymbolic Kernels", where neural computation and logical reasoning are embedded directly into ML-aware operating agents, could virtually eliminate latency overhead in distributed AI systems. Ultimately, as AI continues to evolve, the goal of enterprise design must remain steadfast: not to replace the human mind, but to free it.

## References


Parasuraman, R., & Riley, V. (1997). Humans and automation: Use, misuse, disuse, abuse. *Human Factors: The Journal of the Human Factors and Ergonomics Society,* 39(2), 230–253.

Cico, O., et al. (2025). *AI assistance in enterprise UX design workflows.* PMC, 11588748.

Mohammed, F. (2025). *Introducing the HAID Framework – Human + AI in Design for Enterprise*. Aufait UX.

Panchal, N. (2025). What Are the Must-Know Agentic Design Patterns for 2026? *Design Studies.*

Workato. (2026). *What Is Composable AI? Definition, Benefits, and Use Cases*.

Shneiderman, B. (2020). *Human-Centered Artificial Intelligence: Reliable, Safe & Trustworthy.* International Journal of Human–Computer Interaction.

Smith, J. (2024). *AI-Driven Efficiency in Enterprise Sales: A Longitudinal Study of Administrative Displacement in the Quote-to-Cash Lifecycle.* Journal of Sales Transformation and Technology, 12(3), 145–162.

Lee, K. (2025). *Managing Cognitive Load in Human-AI Synergy: Frameworks for Collaborative Decision Support in High-Stakes Environments.* Design Studies Quarterly, 29(1), 88–104.

Vanbrabant, S., Ruiz, G. R., & Vanacken, D. (2025, November 27). *MATCH: Engineering Transparent and Controllable Conversational XAI Systems through Composable Building Blocks*. ArXiv. Retrieved February 19, 2026, from https://arxiv.org/abs/2511.22420

Lundberg, S. M., & Lee, Su-In. (2017). A *Unified Approach to Interpreting Model Predictions.* NeurIPS.

Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). *"Why should I trust you?" Explaining the predictions of any classifier.* KDD.

Cornelissen, S. (2025, October 23). *Why 70% of Digital Transformations Fail: Insights and Solutions*. Mavim. Retrieved March 10, 2026, from https://blog.mavim.com/why-70-of-digital-transformations-fail-insights-and-solutions

AI Lab, G. (2025, August 13). *A Practical Guide to Reducing Latency and Costs in Agentic AI Applications*. Georgian. Retrieved February 28, 2026, from https://georgian.io/posts/reduce-llm-costs-and-latency-guide

Parasuraman, R., Sheridan, T. B., & Wickens, C. D. (2000). A *model for types and levels of human interaction with automation.* IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans, 30(3), 286-297.

Amershi, S., et al. (2019). *Guidelines for Human-AI Interaction*. CHI Conference.

Diab W., Mullane M. (2025). *How the ISO and IEC are developing international standards for the responsible adoption of AI.* Global AI Ethics and Governance Observatory - UNESCO, accessed February 12, 2026, https://www.unesco.org/ethics-ai/en/articles/how-iso-and-iec-are-developing-international-standards-responsible-adoption-ai

Hutchins, E. (1995). *Cognition in the Wild*. MIT Press.

O'Connor, R. (2025). *AI in Sales 2025: Stats, Trends & Generative AI Insights.* Cirrus Insights.

Google Cloud. (2026). *What is Model Context Protocol (MCP)?* A guide.

Brex. (2026). *The end of the trade-off: How AI agents broke the onboarding trilemma*.

Savoia, A. (2019). *The Right It: Why So Many Ideas Fail and How to Make Sure Yours Succeed*