

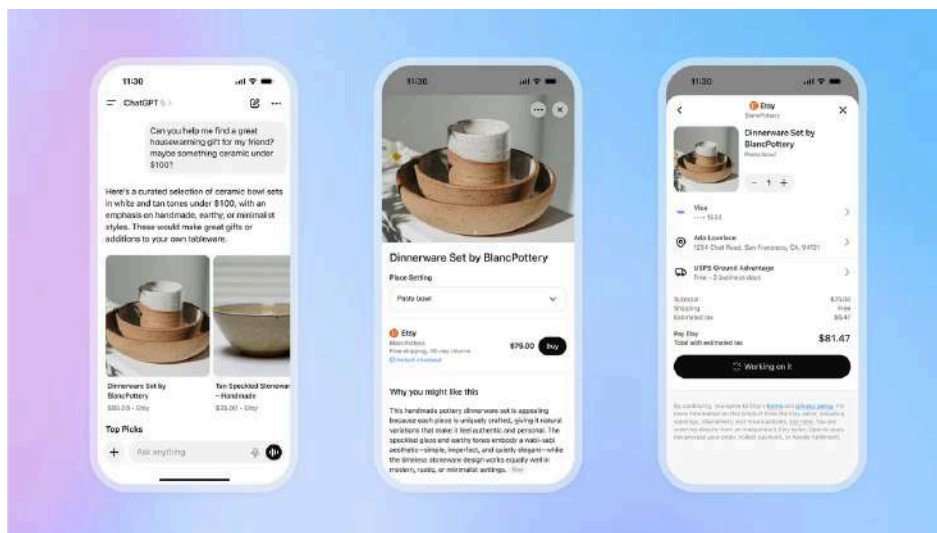
THE EXECUTIVE GUIDE TO **Agentic Commerce**



1.0 Introduction: The Dawn of Agentic Commerce

1.1 Why Now?

In the fourth quarter of 2025, the long-theorized future of e-commerce became a present-day reality. OpenAI, the firm at the forefront of the generative AI revolution, announced its "Instant Checkout" feature, enabling direct, in-conversation purchases within its ChatGPT interface. This was immediately followed by a landmark partnership with Walmart, the world's largest retailer. This collaboration allows Walmart's 270 million weekly customers to browse, discover, and purchase products from the retailer's vast catalog without ever leaving the chat window.



This announcement was not merely an incremental feature release; it was a watershed moment, signaling a fundamental paradigm shift in digital commerce. As Walmart's CEO, Doug McMillon, articulated, the era of the static search bar is ending, giving way to a "native AI experience... that is multi-media, personalized and contextual". This is the arrival of agentic commerce: a new model where AI shifts from a reactive assistant that answers queries to a proactive agent that anticipates needs, makes decisions, and executes complex tasks—like a complete shopping transaction—on behalf of the user. The strategic alignment of the world's leading AI platform with its largest retailer provides an unmistakable signal to the market: the race to own the conversational customer interface has begun, and the stakes are existential.

1.2 Defining the Opportunity: Beyond the Chatbot

To understand the magnitude of this shift, it is crucial to define its core components. Chat-enabled instant checkout refers to a seamless, end-to-end transactional capability within a single conversational interface. It is the ability for a customer to move from initial query ("I need a waterproof jacket for hiking in the Pacific Northwest in November") to a completed purchase, including size selection, shipping, and payment, without ever clicking away to a traditional product grid or checkout page.



This capability is the functional expression of the broader strategic concept of agentic commerce. This is not an evolution of e-commerce; it is a fundamental re-architecting of the customer journey. As analysts at McKinsey describe, agentic commerce is a "seismic shift" where intelligent AI agents act autonomously on behalf of consumers, capable of anticipating needs, personalizing the journey, and automating every step of the process. We are moving from a world where businesses optimize clicks to one where they must earn the trust of algorithms acting as proxies for their customers.

1.3 Quantifying the Business Imperative

The urgency for enterprises to embrace this new paradigm is underscored by the immense, quantifiable failures of the old one. The traditional e-commerce funnel is notoriously inefficient, bleeding revenue at every step. This inefficiency, combined with the proven returns of personalization, creates a powerful and urgent business case for agentic commerce.

- **The Multi-Billion-Dollar Friction Problem:** The global e-commerce cart abandonment rate stands at a staggering 70.19%. This translates into an estimated \$18 billion in lost sales revenue annually for businesses. A primary driver of this abandonment, cited by 22% of

shoppers, is a "long or complicated checkout process." Agentic commerce directly attacks this friction by collapsing the entire multi-page, multi-click funnel into a single, fluid conversation.

- **The Proven ROI of Hyper-Personalization:** The generic, one-size-fits-all nature of traditional e-commerce is another significant headwind to conversion. The data is unequivocal: personalization drives revenue. According to McKinsey, personalization strategies can boost sales by up to 10%, while other studies show that personalized product recommendations can drive 40% more purchases. An AI agent, capable of understanding a user's context, history, and intent, is the ultimate personalization engine, able to tailor recommendations and guide discovery in a way that static web pages cannot.
- **The Future of Customer Interaction:** The strategic importance of this shift is validated by leading industry analysts. Gartner predicts that by 2029, agentic AI will autonomously resolve 80% of common customer service issues, leading to a 30% reduction in operational costs. Forrester forecasts that the global online retail market, the very arena being reshaped by this technology, will reach \$6.8 trillion by 2028. The convergence of these forces—the high-profile market validation from leaders like Walmart, the persistent and costly failures of the existing e-commerce model, and the proven financial upside of true personalization—creates an undeniable strategic imperative. The question for business leaders is no longer if they should invest in agentic commerce, but how they will build the capability to compete and win in this new landscape.

2.0 The New Competitive Landscape: From Novelty to Necessity

The shift to agentic commerce is not a theoretical, future-state concept. It is a present-day reality, with pioneers across retail, food service, and B2B sectors already deploying conversational solutions to solve critical business problems. An examination of these early adopters reveals a clear evolutionary path, from simple channel experiments to deeply integrated, proprietary experience engines that are becoming a core source of competitive advantage. This evidence of market maturity underscores the urgency for every enterprise to develop its own strategy.

2.1 Pioneers in Conversational & Agentic Commerce

The following table provides an overview of how leading brands have leveraged conversational and agentic commerce, highlighting a clear progression in sophistication and strategic intent.

Company	Core Functionality	Platform / Channel
---------	--------------------	--------------------

Walmart	Full Product Discovery & Instant Checkout	ChatGPT (Third-Party Integration)
Amazon	AI-Powered Conversational Shopping Assistant	"Rufus" (Proprietary In-App)
Domino's Pizza	Re-ordering (Easy Order) & Order Tracking	"AnyWare" Platform (Alexa, Text, Watch)
Wendy's	Full Order Taking (Voice)	"FreshAI" (Proprietary Drive-Thru)
Jet's Pizza / White Castle	Full Order Taking (Voice)	Third-Party Voice AI (Drive-Thru/Phone)
1-800-Flowers	Guided Full Order & Checkout	Facebook Messenger Bot
Sephora	Product Tutorials & Recommendations	Kik & Messenger Bots
Taco Bell	Menu Questions & Full Order Taking	"TacoBot" (Slack)
Pizza Hut	Re-ordering & Deals	Messenger & Twitter Bots
Starbucks	Voice Re-ordering & New Orders	"My Starbucks Barista" (In-App & Alexa)
Staples	Customer Support & Future Order Integration	"Easy System" (App, Messenger)

2.2 Analysis: The Evolutionary Path of Agentic Commerce

The examples above are not random; they trace a clear trajectory of increasing strategic importance and technical sophistication, a journey that every enterprise must now consider navigating.

- **Phase 1:** The Channel Experiment (c. 2016). The first wave of conversational commerce was about establishing a presence on emerging third-party messaging platforms. Initiatives like Taco Bell's "TacoBot" on Slack and 1-800-Flowers on Facebook Messenger were primarily

experiments in meeting customers where they were. The functionality was often limited to simple, guided flows or basic re-orders, like those offered by Pizza Hut on Messenger and Twitter. The primary goal was channel presence and brand novelty, not deep operational integration.

- **Phase 2:** The Efficiency Play. The next phase focused on leveraging conversational interfaces to automate high-frequency, low-complexity tasks to drive operational efficiency and customer loyalty. Domino's "AnyWare" platform is the canonical example of this stage. By enabling customers to place their saved "Easy Order" via text, Amazon Alexa, or an Apple Watch, Domino's streamlined the re-ordering process, making it the path of least resistance and cementing customer habit. The strategic goal shifted from novelty to utility and retention.
- **Phase 3:** The Experience Engine. The current and most critical phase involves the creation of proprietary, deeply integrated AI agents that are core to the customer experience. Amazon's "Rufus" assistant, built directly into its mobile app, is not just a chatbot; it is a sophisticated shopping guide that can handle complex, open-ended queries, compare products, and draw insights from customer reviews. Similarly, Wendy's "FreshAI" drive-thru system is a proprietary voice agent designed to handle the full complexity of a new order, with the explicit goals of improving order accuracy, increasing average check size through intelligent upselling, and boosting profit margins. In this phase, the AI agent is not an ancillary channel; it is a primary engine of revenue and differentiation.

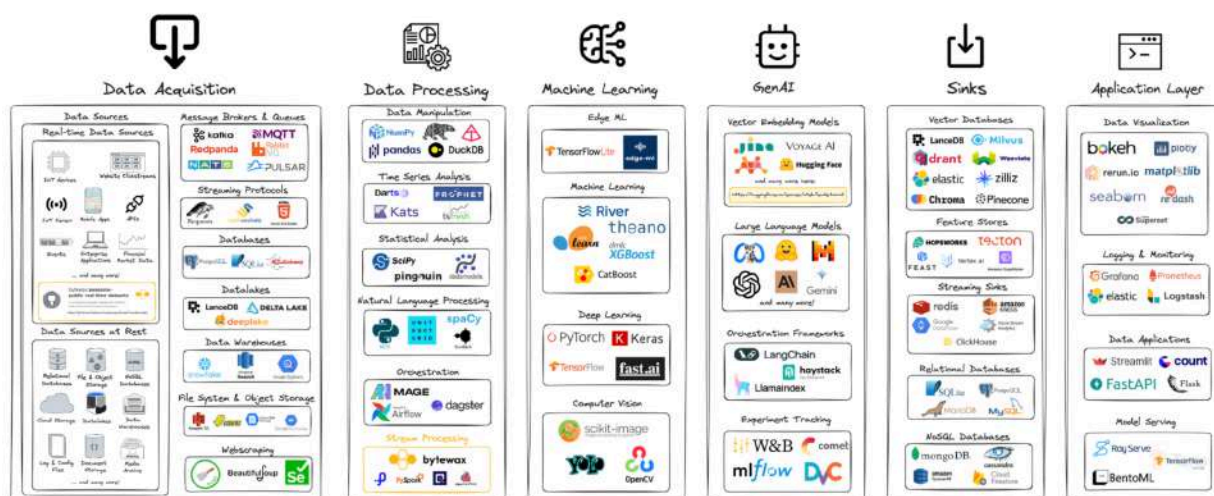
This evolutionary path reveals a critical pattern. As organizations progress from simply experimenting with conversational commerce to treating it as a core strategic function, they invariably conclude that they must own and control the underlying platform. The early reliance on third-party messaging apps like Kik and Messenger gives way to the development of proprietary, in-house systems like Rufus and FreshAI. This strategic migration from renting to owning the customer interface is a direct response to the inherent limitations and risks of building a core business function on a platform one does not control. This leads directly to the central strategic question every business leader must now face.

3.0 The Central Question: The "Easy" Path vs. The Strategic Path

With the imperative to act now clearly established, enterprise leaders are faced with a critical decision: integrate with a third-party AI platform, or build a proprietary, in-house solution. The first path offers the allure of speed, while the second presents the opportunity for sustainable competitive advantage. For any serious enterprise, particularly those in regulated or data-sensitive industries, a thorough

analysis reveals that the "easy" path is fraught with unacceptable risks, making the strategic path of building in-house the only viable option for long-term success.

3.1 The Siren Song of Third-Party Integration



The appeal of leveraging a powerful, pre-trained Large Language Model (LLM) via a third-party API, such as those offered by OpenAI, is undeniable. It promises rapid development, lower upfront investment in model training, and immediate access to state-of-the-art AI capabilities. However, this path requires a business to build its core customer experience on a foundation it neither owns nor controls, introducing a set of profound and often enterprise-ending risks.

- Risk 1: Catastrophic Data Governance & Compliance Failure.** The moment a customer's conversational data—containing personally identifiable information (PII), product preferences, and payment details—is sent to a third-party API, it leaves the enterprise's governance boundary. This creates a massive compliance liability. Strict regulations like the EU's General Data Protection Regulation (GDPR), the California Consumer Privacy Act (CCPA), the Health Insurance Portability and Accountability Act (HIPAA), and the Payment Card Industry Data Security Standard (PCI DSS) all mandate rigorous data protection, auditability, and control. The infamous Facebook API data leak, which exposed the data of millions, stands as a stark reminder of the consequences of insufficient governance over third-party data access. For any enterprise in finance, healthcare, or any sector that processes payments, outsourcing the raw conversational interface to a third party is a compliance non-starter that invites multi-million-dollar fines and irreparable reputational damage.

- **Risk 2: Crippling Vendor Lock-In.** Building a core business function on an external platform creates a dangerous strategic dependency. The business becomes beholden to the vendor's roadmap, pricing model, and terms of service. What happens when the API provider decides to 10x their pricing, deprecate the model version the system is built upon, or change their data usage policies? The entire chat-to-checkout channel is held hostage. The abrupt 2022 shutdown of Google Cloud's IoT Core service, which forced customers into costly and frantic emergency migrations, illustrates the very real danger of this dependency.
- **Risk 3: Loss of Control and Competitive Differentiation.** Third-party models are, by their nature, generic. While they can be fine-tuned, achieving deep customization to perfectly reflect a unique brand voice, enforce complex and dynamic business rules, or query proprietary, real-time data sources (like live inventory or customer-specific pricing) is exceptionally difficult. This inevitably leads to a commoditized, "good enough" customer experience that fails to create a meaningful competitive advantage.
- **Risk 4: Existential Threat of Competitive Data Leakage.** This is the most insidious risk of all. An enterprise's proprietary data—its sales trends, customer query patterns, and conversion funnels—is its most valuable strategic asset. When this data is sent to a third-party AI provider, it is often used to train and improve their next-generation models. The enterprise is, in effect, paying to educate the vendor's AI. That same, newly improved model can then be licensed to a direct competitor, who now benefits from the insights gleaned from the first company's customers. This is not merely a data security risk; it is the strategic folly of actively funding the research and development of one's competition.

3.2 The Strategic Imperative of Building In-House

Given the profound risks of the third-party approach, building a proprietary, in-house agentic commerce solution emerges as the only strategic path for enterprises seeking to create a durable competitive advantage. This approach allows a company to achieve three critical business outcomes that are impossible with an outsourced model:

- **Absolute Data Control & Security:** By keeping all sensitive customer conversations, PII, and transaction data within the company's own secure perimeter, an in-house solution ensures full compliance with data regulations and builds the foundation of customer trust essential for this new, more intimate form of commerce.

- **True Differentiation and a Bespoke Experience:** An in-house agent can be meticulously crafted to embody the company's brand, understand the nuances of its unique product catalog, and execute complex business logic that reflects its specific go-to-market strategy. This allows for the creation of a truly differentiated, premium customer experience that cannot be replicated by competitors using off-the-shelf tools.
- **A Defensible Competitive Moat:** A proprietary agent, trained on a company's unique data and deeply integrated with its core operational systems (ERP, CRM, logistics), becomes a core piece of intellectual property. It is a strategic asset that grows more valuable with every customer interaction, creating a powerful, self-reinforcing competitive moat.

While the strategic benefits are clear, the decision to build in-house is not trivial. The technical complexity of executing this strategy is immense, and it is here that the majority of internal AI projects falter, long before they deliver any business value.

4.0 The Three Pillars of Failure: Why In-House AI Projects Stall

The strategic case for building a proprietary agentic commerce solution is compelling. However, the path is littered with failed projects, budget overruns, and stalled initiatives. The reasons for this are consistent and predictable. Enterprises that attempt to build these complex systems from the ground up invariably run into three fundamental, platform-level challenges that consume vast resources and ultimately derail their efforts. These projects fail not because the AI models are flawed, but because the organization cannot provide the secure, integrated, and scalable foundation required to run those models in a production environment.

4.1 Challenge 1: The Data Governance & Security Nightmare

The foundational challenge of any in-house agentic commerce project is a paradox of security: how can an enterprise grant a powerful, probabilistic AI model real-time, read/write access to its most sensitive "crown jewel" data systems without that data ever leaving its secure, governed perimeter?

An effective chat-to-checkout agent must be able to query a live inventory database, access customer PII and order history from a CRM, pull dynamic pricing from an ERP, and interact with a payment gateway. These systems are the operational heart of the business. According to Deloitte, while generative AI strategies demand access to massive, diverse data sets, business leaders are deeply concerned about data privacy, security, data sovereignty, and governance. Most legacy enterprise data environments are built on deterministic, siloed architectures that are fundamentally ill-equipped to provide the secure, real-time, and flexible data access that modern AI systems require. Before a single

line of the agent's business logic can be written, a complex security and governance framework must be engineered to create a "virtual air-gap" around the AI, allowing it to access the data it needs without exposing the enterprise to catastrophic risk. For most organizations, building this bespoke security layer is a project in itself, often taking years and consuming the majority of the initiative's budget.

4.2 Challenge 2: The "Plumbing" & Orchestration Mess

A chat-to-checkout agent is not a monolithic AI application. It is a complex, distributed system that requires the real-time orchestration of numerous, disparate components. The "plumbing" required to connect these components is the second pillar of failure where promising AI projects go to die.

Consider a simple customer request: "Do you have the blue running shoes from my last order in a size 10, can they ship to me by Friday, and can I pay with my rewards points?" To answer this, the AI agent must orchestrate a sequence of near-instantaneous calls to multiple systems:

1. Query the CRM to retrieve the customer's order history.
2. Query the Product Information Management (PIM) system to get the SKU for the blue version of the shoe.
3. Query the inventory database or ERP to check stock levels for that SKU in a size 10.
4. Query the logistics and shipping API to calculate a delivery date.
5. Query the loyalty platform to check the customer's points balance.
6. Synthesize all of this information into a coherent, natural language response.
7. Finally, upon confirmation, orchestrate calls to the order management system and payment gateway.

This intricate dance of data must happen in milliseconds. However, in most enterprises, these critical data sources exist in isolated silos, often within legacy systems that do not easily communicate with one another. The data engineering and DevOps effort required to build and maintain this real-time orchestration layer from scratch is monumental. It involves managing dozens of different APIs, authentication methods, data formats, and network latencies, a complexity that quickly overwhelms internal teams and stalls progress indefinitely.

4.3 Challenge 3: The MLOps & Scalability Mountain

Even if an enterprise can solve the security and orchestration challenges, it faces a third, formidable obstacle: the mountain of non-differentiating but mission-critical infrastructure work required to run

AI models in production at scale. This discipline, known as MLOps (Machine Learning Operations), is a notorious source of complexity and cost.

Successfully deploying and managing an agentic commerce system requires a vast array of MLOps capabilities: provisioning and managing clusters of specialized GPU hardware; building automated pipelines for model training, versioning, and deployment; implementing robust systems for logging, monitoring, and auditing every request; and engineering a system that can autoscale from a handful of concurrent users to millions during a peak sales event like Black Friday. Building this infrastructure from the ground up is a cripplingly expensive and time-consuming endeavor. It requires hiring teams of scarce and highly paid MLOps specialists, data engineers, and cloud infrastructure experts. The average time to build a proficient internal AI team and the necessary infrastructure can range from six months to two years, with costs for specialized talent, cloud services, and GPU hardware quickly running into the millions. This is the undifferentiated heavy lifting that provides no direct competitive advantage but is absolutely essential for a production-grade system. It is a resource drain that distracts the organization's best technical talent from the actual goal: building a great customer-facing product.

5.0 The Solution: The Foundational Power of an AI Operating System

The three pillars of failure—insurmountable security challenges, complex data orchestration, and the crushing weight of MLOps—make it clear that attempting to build a proprietary agentic commerce solution from scratch is an untenable strategy for all but a handful of hyperscale technology companies. This reality does not, however, invalidate the strategic imperative to build in-house. Instead, it points to the need for a new foundational layer in the enterprise technology stack: an AI Operating System.

5.1 A New Architectural Layer: The AI Operating System

Just as a computer's operating system (OS) handles the complex, low-level tasks of memory management, process scheduling, and hardware communication, allowing developers to focus on writing applications, an AI Operating System handles the complex, undifferentiated "plumbing" of enterprise AI. It provides a unified, production-grade platform that solves the foundational challenges of security, orchestration, and scalability, allowing an enterprise's AI and data science teams to focus exclusively on building the high-value, business-differentiating logic of their conversational agents.

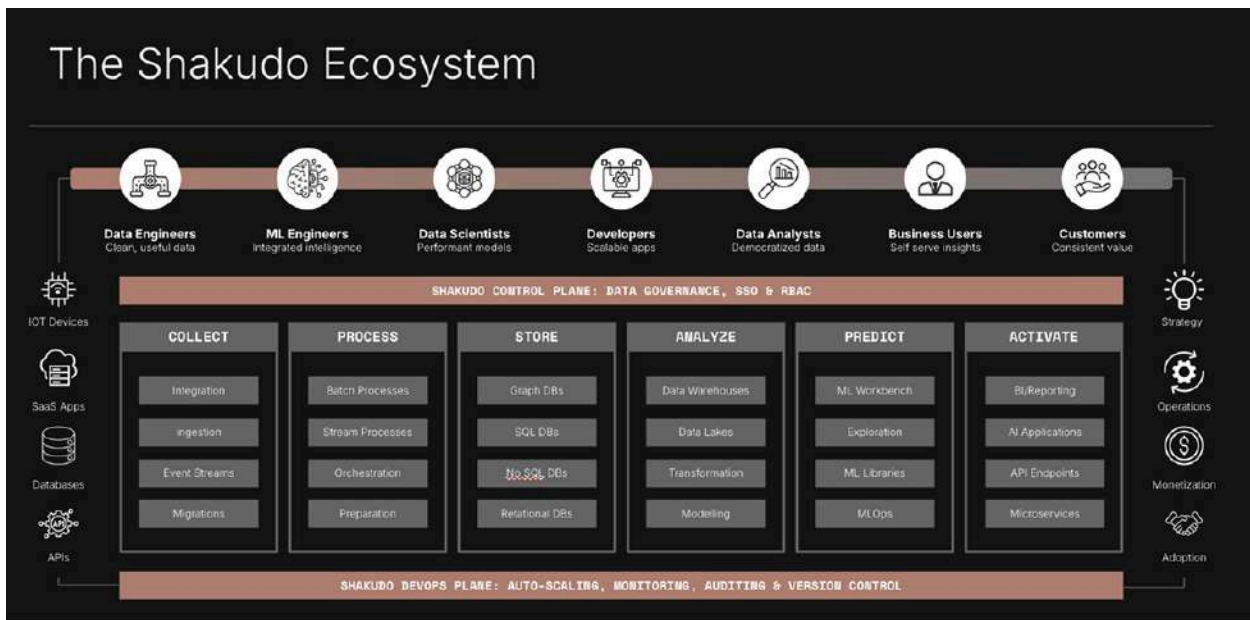
This new architectural layer provides the central control plane for all data and AI initiatives. It abstracts away the undifferentiated heavy lifting detailed in the previous section, transforming the process of

building a secure, scalable, and integrated AI application from a multi-year infrastructure project into a focused software development effort.

5.2 Shakudo: The Enterprise-Native AI Operating System

Shakudo is the platform built from the ground up to be this exact operating system for enterprise data and AI. It is a comprehensive, integrated platform that provides all the foundational capabilities needed to build, deploy, and scale sophisticated AI applications like a chat-enabled instant checkout agent.

The most critical architectural feature of the Shakudo platform is that it deploys entirely inside a company's own infrastructure—whether in a private cloud VPC or on-premise data center. This single design principle is the key that unlocks the solution to the primary challenges that derail in-house AI projects. By operating within the enterprise's existing security perimeter, Shakudo enables organizations to maintain absolute control over their data and governance while leveraging the full power of modern AI.



5.3 From Challenge to Solution with Shakudo

The Shakudo platform is built on three core pillars, each designed to directly solve one of the three pillars of failure for in-house AI development. The mapping is direct and unambiguous, providing a clear path for enterprises to overcome the obstacles to building their own agentic commerce solutions.

In-House Challenge	The Shakudo Solution Pillar	Business Outcome
1. Data Governance & Security Nightmare	Absolute Control & Governance	Securely connect LLMs to proprietary "crown jewel" data inside your own VPC. Sensitive data never leaves your governance boundary, ensuring full compliance (GDPR, CCPA, HIPAA, PCI).
2. "Plumbing" & Orchestration Mess	Tool-Agnostic Orchestration	Seamlessly connect all the necessary systems—LLMs (any model), vector DBs, CRMs, ERPs, payment gateways—from a single, unified control plane. Eliminate integration complexity and vendor lock-in.
3. MLOps & Scalability Mountain	Production-Grade Scalability	Automate the entire MLOps/DevOps stack, from GPU cluster management to autoscaling. Reduce deployment time from months to weeks and scale instantly from 10 to 10 million users.

By providing this enterprise-native operating system, Shakudo fundamentally changes the calculus of the "build vs. integrate" decision. It eliminates the prohibitive complexity and risk of building the foundational platform from scratch, enabling enterprises to pursue the strategic path of owning their AI-driven customer experience with confidence and speed.

6.0 Focus: High-Stakes Industries

While the strategic imperative to build a proprietary agentic commerce solution applies to all consumer-facing enterprises, for certain high-stakes industries, it is not merely a competitive advantage—it is a regulatory and operational necessity. In sectors defined by sensitive data, complex products, and strict compliance regimes, the "build on a secure foundation" approach is the only viable path forward.

- Banking & Financial Services:** The financial sector is rapidly adopting conversational AI for use cases like new customer onboarding, loan applications, and 24/7 customer support. These interactions invariably involve the exchange of highly sensitive financial data and PII. The non-negotiable requirements of regulations like PCI DSS and GDPR make the use of third-party APIs, which send this data outside the bank's governance boundary, an unacceptable compliance risk. An enterprise-native platform like Shakudo, which operates

entirely within the bank's secure infrastructure, provides the foundational control and auditability required to deploy these services safely.

- **Healthcare:** The applications for conversational AI in healthcare are transformative, ranging from compliant appointment booking and prescription re-ordering to patient triage and post-care follow-up. Every one of these use cases involves Protected Health Information (PHI), which is governed by the stringent privacy and security rules of HIPAA. Absolute control over where this data resides and who can access it is paramount. Building these solutions on an enterprise-native AI operating system is essential to ensure compliance and protect patient privacy.
- **B2B Manufacturing & Distribution:** The B2B commerce landscape is characterized by complexity. Buyers are not making simple, one-off purchases; they are re-ordering specific SKUs, requesting real-time quotes based on volume and contract terms, and checking live inventory levels for mission-critical parts. An effective B2B conversational agent must have deep, real-time, and permission-aware integration with core back-end systems like ERPs, CRMs, and inventory management platforms. This level of deep, secure integration is impossible to achieve with generic, external AI models but is a core function of a tool-agnostic orchestration platform like Shakudo.

In these industries, the decision is not a choice between speed and control. The regulatory and operational realities mandate control. The role of a platform like Shakudo is to provide that control without sacrificing the speed and agility needed to compete.

7.0 Conclusion: From Cost Center to Profit Center

The commercial landscape is at an inflection point. The emergence of true agentic commerce, validated by market leaders like Walmart, has redefined the future of the customer experience. The static, friction-filled e-commerce funnel is being replaced by a dynamic, personalized, and seamless conversational interface. For enterprise leaders, this is not a trend to be monitored; it is a fundamental shift that demands a strategic response.

This guide has laid out the stark choice facing every organization. The "easy" path of integrating with third-party AI platforms is a strategic trap. It offers the illusion of speed while exposing the enterprise to unacceptable risks in data governance, vendor lock-in, and, most critically, the leakage of proprietary data that can be used to arm competitors. It is a path that leads to commoditization and strategic vulnerability.

The only path to sustainable competitive advantage is to build and own a proprietary agentic commerce solution. This is the only way to ensure absolute data security, create a truly differentiated brand experience, and build a defensible competitive moat. However, the technical complexity of this path—the nightmare of data governance, the mess of systems orchestration, and the mountain of MLOps—has historically made it an untenable option for most. This is why so many internal AI projects become high-risk, complex cost centers that fail to deliver value.

This is where the paradigm of an AI Operating System changes the equation. A foundational platform like Shakudo, which deploys natively within an enterprise's own secure infrastructure, solves the undifferentiated but mission-critical challenges of security, orchestration, and scalability. It provides the solid foundation upon which an organization's data scientists and developers can build what truly matters: the unique, intelligent, and brand-aligned agent that will define their customer relationships for the next decade.

By adopting this platform-centric approach, an enterprise transforms its AI initiative. It is no longer a risky, multi-year infrastructure project but a focused, agile development effort. It is a shift from a complex cost center to a powerful, secure, and scalable profit center that will drive revenue and customer loyalty long into the future.

ABOUT SHAKUDO

Shakudo is the operating system for AI, existing completely within your private infrastructure and guaranteeing absolute control, data governance, and faster time to market than ever. Like an operating system, Shakudo streamlines AI adoption through a tool-agnostic orchestration approach using the best-of-breed technologies that eliminates complex DevOps overhead, vendor lock-in, and security vulnerabilities. Organizations across critical sectors—including finance, energy, and defense—choose Shakudo for its guaranteed time-to-value.

Shakudo delivers ultimate scalability and speed, allowing teams to focus on driving business outcomes.

Shakudo: Intelligence without constraint. Find out more at shakudo.io.

