Agentic Commerce

First-mover advantage: Everything you need to know and do now

shopware

Table of contents

Chapter 1

Laying the groundwork: Understanding agentic commerce

shopware

Introduction

The language of ecommerce has become densely populated with frameworks rich in ambition: omnichannel, composable, headless, modular, spatial, experiential—each attempting to describe a future that continues to accelerate. Now, a new term enters the scene: **agentic commerce**.

There's no shortage of noise surrounding agentic commerce. Across the industry, definitions vary, expectations diverge, and some critics are already dismissing it as the next over-hyped buzzword. But peel back the jargon, and something more consequential emerges. If the right conditions are met (data infrastructure, model intelligence, ethical guardrails) agentic AI has the potential to fundamentally reconfigure the architecture and logic of digital commerce.

From interface to intent: Defining agentic commerce

At its core, agentic commerce transfers execution (but not authority) from the end user and/or the merchant to autonomous software agents acting independently **on their behalf**. Where traditional ecommerce requires constant human input, agentic systems are designed to interpret intent and act upon it. These agents could manage product discovery, personalize interactions, execute transactions, and even handle post-purchase engagement, liberating users from decision fatigue and reducing complexity across the value chain.

For merchants, this shift unlocks a new operational paradigm. Instead of investing time and resources into repetitive or reactive tasks, businesses can delegate a wide range of functions to intelligent agents: from planning campaigns and ideating product bundles to managing inventories, analyzing performance data, and generating on-brand content at scale. Agentic systems empower teams to focus on strategic growth while automating the orchestration of everyday execution.

The transformation is both architectural and behavioral. But more than that, it signals a deeper structural redefinition of what digital commerce is. Just as the internet collapsed the boundaries of time and space, agentic commerce collapses the divide between interaction and execution. Agentic AI doesn't just automate – it delegates cognitively, enabling machines to complete tasks and make informed decisions.

If realized responsibly, the implications for business models, user experience, and operational design will be huge. Agentic AI is not a new layer atop ecommerce. It's a re-ordering of its logic. And in that sense, it can be the most profound turning point since the invention of the web itself.

This white paper explores agentic commerce as both a strategic imperative and a practical opportunity. It examines why this shift from human-led interaction to autonomous intent execution is happening now, what it means for the future of B2B and B2C commerce, and how businesses can begin preparing today.

We also take a critical look at the current limitations of agentic systems, from infrastructural and algorithmic challenges to questions of trust, oversight, and energy use, ensuring the conversation remains both ambitious and grounded. And because Shopware is deeply invested in shaping this future, we also share our vision: one built on open systems, scalable architecture, and merchant-first innovation that empowers real people to achieve more with less friction.

Why now? The convergence of architecture, intelligence, and engagement

Agentic commerce is not emerging in a vacuum. It's the result of several long-developing forces that have reached critical maturity at the same time: technological, infrastructural, and behavioral. While the underlying ideas have existed in academic and enterprise contexts for years, only recently have the conditions aligned to make autonomous, commercially viable software agents possible.

The enabling conditions of the present moment

Three converging dynamics have made agentic commerce both technically viable and increasingly necessary:

- **The API economy** has transformed the architectural landscape of digital commerce. Advanced platforms now expose their core services (inventory, checkout, pricing engines, loyalty systems, customer data, etc.) through APIs. This allows agents to operate in ecosystems without friction, orchestrating complex workflows without the overhead of direct human interaction.
- **The maturation of AI** has removed key limitations on agent autonomy. Large language models (LLMs), advanced deep learning architectures, and reinforcement learning techniques have evolved to the point where software agents can understand nuance, adapt to context, and optimize outcomes over time. These systems are no longer merely reactive. They can reason, communicate, and make trade-offs in open-ended environments. Crucially, they now interact with structured and unstructured data alike, bridging the gap between decision-making and execution.
- **Changing user behavior** has created a pressing demand for automation that is not just efficient, but empathetic. Consumers, particularly digital natives, have grown increasingly fatigued by choice overload, generic personalization, and fragmented customer journeys. Their expectations have shifted toward seamless experiences where intention is understood and acted upon without the need for repeated inputs. The same pattern holds true in B2B as well: buyers now expect supplier interactions to be as smooth, intelligent, and responsive as their consumer apps.

From readiness to inevitability

Together, these forces have created an environment where autonomous agents can realistically perform meaningful commercial tasks, often with greater consistency, scale, and speed than traditional rule-based systems or manual workflows. They can manage replenishment, configure dynamic pricing, initiate transactions, or even act as decision-support layers for human teams.

Importantly, this shift is no longer speculative. As Gartner noted in its Future of Sales 2024 report: **by 2027, one in four B2B transactions will be initiated by AI agents operating independently**. (source) This prediction reflects a broader recognition that the logic of commerce is changing: from human-triggered action to goal-oriented delegation, from interface to intent. Agentic commerce has moved from possibility to inevitability. What remains is the question of readiness.

The forces behind autonomy: Generative AI and machine learning

Autonomy in agentic commerce is not possible without intelligence. The ability for software agents to act independently, make context-aware decisions, and communicate effectively depends on two foundational technologies: **machine learning** and **generative AI**. Together, these technologies form the cognitive substrate of agentic systems, enabling them to reason, adapt, and interact at scale.

Learning from data: The role of machine learning

Machine learning provides the adaptive backbone of agentic commerce. Models, particularly those based on deep learning architectures, allow agents to process, evaluate, and learn from vast streams of data. These include user behavior, purchase history, product performance, inventory fluctuations, and real-time environmental inputs such as competitor pricing or regional demand signals.

Through **supervised learning**, **collaborative filtering**, and **reinforcement learning**, agents can optimize their actions toward specific goals, whether maximizing conversion, reducing churn, or personalizing product recommendations. Over time, agents evolve through feedback loops, improving their decision quality and contextual relevance with every interaction.

Critically, this learning is not static. Agentic systems continuously adapt to shifting inputs, enabling them to operate reliably even in unpredictable or high-variance environments. This makes them far more flexible than rule-based automation, which often breaks under complexity.

Communicating with users: The generative layer

If machine learning gives agents the capacity to decide, **generative AI gives them the capacity to communicate**. Transformer-based large language models (LLMs) such as GPT, Claude, and Gemini enable agents to synthesize human-like language in real time. This is essential not only for customer-facing interactions (think chatbot responses, product descriptions, or negotiation scripts) but also for internal processes such as summarizing analytics or translating data into insights for teams.

Generative AI allows agents to act as interpreters between structured systems and unstructured environments. They can turn intent into action, data into dialogue, and user signals into tailored messaging. In this sense, generative capabilities are not cosmetic. They are core to making autonomy intelligible and accessible within the commerce stack.

Synergy at scale

The true power of agentic commerce lies in the synergy between adaptive learning and generative fluency. Together, these capabilities enable agents to **anticipate needs, personalize experiences, and explain their reasoning**. This moves agentic commerce beyond rigid workflows into a space of proactive orchestration where decisions are contextual, communication is natural, and outcomes are optimized in real time. It is this blend of continuous learning and expressive capability that allows agentic systems to operate effectively within complex, multi-layered ecommerce environments.

Adaptive by design: Feedback loops and domain expertise

Learning in motion through feedback loops

Intelligence without feedback is brittle. Agentic systems depend on continuous learning, and at the heart of this adaptability lies the feedback loop. These loops capture user behavior, transaction outcomes, and key success metrics, feeding this data back into the agent's learning process. They are made possible through mechanisms such as interaction signals, outcome-based evaluation, and reinforcement learning, where agents iteratively refine their actions based on rewards, penalties, and observed performance over time.

This constant calibration transforms agents from static rule-followers into dynamic decision-makers. Rather than executing predefined workflows, agents learn to respond to changing user preferences, business goals, and environmental inputs, improving over time with each iteration. Feedback is a structural necessity for autonomous systems to remain effective in complex, high-variance environments like ecommerce.

Domain expertise as a constraint and enabler

Learning alone is not enough. For agentic systems to scale across business functions and remain contextually appropriate, they must also possess **domain expertise**: the vertical-specific knowledge that shapes what "good" decisions look like in a given setting. This expertise is embedded into task-specific models or encoded in knowledge graphs that reflect industry logic such as regulatory constraints in pharmaceuticals, seasonality patterns in fashion, bundling strategies in electronics, or logistics thresholds in manufacturing.

Domain-aware agents can operate autonomously within these defined boundaries while contributing to a broader agentic ecosystem. They don't just optimize for outcomes. They optimize for outcomes that make sense within the operational, legal, and strategic realities of a business. This ensures decisions remain both performant and grounded, aligning machine logic with human priorities.

From isolated intelligence to orchestrated autonomy

As ecommerce systems evolve from rule-based automation to agentic intelligence, the interplay between feedback loops and domain expertise becomes central to their effectiveness. These mechanisms ensure that agents do not simply execute tasks, but learn from outcomes, refine their strategies, and adapt to changing conditions.

Is this merely a speculative future vision? Far from it. Many of its components are already visible in production environments: AI-driven recommendation engines that learn from user signals, dynamic pricing systems that respond to demand fluctuations in real time, and intelligent search platforms tailored to industry-specific merchandising logic.

What remains emergent is the orchestration of these elements into cohesive, fully autonomous agents capable of managing commercial workflows end to end. The transition is underway. And it is not defined by a single breakthrough, but by the steady convergence of data infrastructure, machine learning, and composable system design.

Chapter 2

Operationalizing agentic commerce: Strategy, scenarios, and business value

shopware

The backbone of autonomy: Open source, APIs, and MCP

Agentic commerce cannot thrive in closed systems. Autonomous agents require infrastructure that is flexible, modular, and accessible, capable of exposing functionality and data in ways that support continuous, intelligent orchestration. This is why **open source** and **API-first architectures** are critical. Together, they form the operational substrate that enables agents not only to function, but to evolve.

Exposing logic through APIs

In traditional monolithic systems, business logic is often tightly coupled, undocumented, or buried deep within proprietary layers. This presents a structural barrier for autonomous agents, which must be able to initiate actions without human intervention: querying a product database, adjusting pricing, launching a campaign. An **API-first approach** solves this by exposing each system component through modular, well-documented endpoints. From checkout workflows to inventory data and user profiles, agents gain real-time access to the services they need to make informed, autonomous decisions.

This architectural philosophy also decouples development cycles, allowing systems to evolve without compromising interoperability. Agents can be integrated incrementally, tested in isolation, and scaled without reengineering the entire stack. It is this decoupling between layers, functions, and services that gives agentic commerce its operational agility.

Open source: Enabling innovation

While APIs provide the interface, open source provides the ecosystem. Open source software gives developers full visibility into the systems that agents rely on. It enables inspection, customization, and optimization, allowing teams to adapt core logic to specific business needs or industry requirements. Just as importantly, it invites contribution. Innovations in agent orchestration, data schema design, or decision intelligence can propagate through a global developer community rather than remain siloed within vendors.

Agentic systems are by nature iterative and adaptive. Closed architectures inhibit their growth. Open systems support the kind of composability that agents need to move beyond simple automation and toward contextual, domain-specific intelligence. The result is an infrastructure that is not just agent-ready, but **agent-native**: built from the ground up to accommodate autonomous systems, plug-in intelligence, and real-time orchestration.

Model Context Protocol (MCP)

For AI agents to operate autonomously, they must do more than just access data—they need to understand it in context. The **Model Context Protocol (MCP)** is a new specification designed to make that possible. It provides a structured, standardized way for agents to retain and apply contextual knowledge across different interactions, workflows, and systems. In the realm of agentic commerce, this has practical implications. An AI agent, for example, can remember a customer's preferences while querying a live inventory API, or reference marketing performance data while writing a product description. It doesn't lose track of what matters, even as it shifts between tools.

What sets MCP apart is its neutrality. Originally developed by Anthropic, the protocol has been rapidly adopted and endorsed by leading players across the AI ecosystem, including OpenAI, Google, Microsoft, and Meta, who now actively use it in their own systems. This broad acceptance speaks to MCP's value as a vendor-agnostic standard. By decoupling context management from any single application or provider, MCP keeps agentic systems **interoperable, portable, and continuously aware**. When paired with open source and API-first design, it becomes the connective layer that enables agents to coordinate complex tasks across domains, with accuracy and ease.

Making it real:
How agentic commerce works in practice

Moving beyond abstraction: what does agentic commerce look like in the wild? Not as a theoretical construct or technical blueprint, but as a living, operational reality embedded in real businesses. It's one thing to define agentic systems in terms of autonomy, feedback loops, or composable infrastructure. It's another to observe how they function in unpredictable, high-variance environments where customers, suppliers, and markets collide in real time.

This chapter explores early signals from the field where intelligent agents are already performing discrete tasks, coordinating decisions, and delivering tangible value. These examples illustrate not just what's possible, but what's already happening at the intersection of autonomy, adaptability, and commercial intent.

B2C: The self-optimizing grocery basket

In the consumer space, one of the clearest expressions of agentic commerce is emerging in automated replenishment and lifestyle-adaptive ordering. Imagine a household subscribed to a weekly grocery delivery service. An agent, trained on family dietary restrictions, past purchasing behavior, seasonal availability and promotions, autonomously assembles a weekly basket. It adjusts not only to product availability but to contextual signals: calendar entries noting visiting guests, updated allergy data, or a change in health goals synced from a fitness app.

The agent applies loyalty points, avoids out-of-stock items, and ensures freshness windows are respected. It doesn't simply react to requests, it **anticipates need and acts accordingly**, reducing friction while enhancing perceived intelligence. Cart abandonment vanishes. Personalization occurs without surveillance. And the consumer experiences not choice overload, but intelligent curation.

B2B: Procurement-as-a-Service

In the B2B domain, agentic commerce offers even more immediate efficiencies. Consider a mid-sized construction firm that must source certified, sustainable building materials from multiple regional suppliers—each with varying pricing models, lead times, and documentation requirements. Historically, this procurement task would involve manual research, compliance checks, quote comparison, and negotiation.

Enter a domain-trained software agent. It understands the company's procurement policies, ESG commitments, and project timelines. It autonomously scans supplier APIs and digital catalogues, filters for materials that meet specified certifications, evaluates historical reliability data, negotiates within set pricing thresholds, and places orders—all while updating internal systems with delivery schedules and cost forecasts.

Here, the agent isn't replacing strategic thinking. It's **removing the operational drag** that prevents procurement teams from focusing on supplier relationships, quality assurance, and long-term planning. It becomes a force multiplier: one that handles complexity at speed and scale.

These examples reflect just the beginning. In both B2B and B2C settings, agentic commerce is not defined by a monolithic rollout but by the **gradual layering of intelligence**, from narrow task automation to multi-domain orchestration. Each agent deployed reduces cognitive load, compresses cycle time, and adds another building block to an infrastructure where decisions are increasingly distributed and data-driven.

This is what agentic commerce looks like in reality: less like a revolution, and more like a shift in gravity, pulling every commercial interaction toward greater autonomy, precision, and adaptability.

Differentiation: Competing at the level of meaning

For brands, agentic commerce demands a new form of differentiation: one that serves both human engagement and machine interpretation. As AI agents increasingly mediate the path to purchase, bypassing traditional browsing, compressing funnels, and optimizing for utility, the risk of commoditization rises sharply. Products become data points, and brands risk becoming invisible.

In this new landscape, where autonomous systems execute on behalf of users, businesses can no longer rely on visibility alone. They must instead compete on distinctiveness and deliver experiences that are both emotionally resonant and structurally legible to machines. The challenge for brands is no longer just how to attract human attention, but how to **stand out in the decision models of intelligent agents**. That means building systems that reflect what makes a business truly different, and making sure both people and machines can see it.

The commoditization problem

Autonomous agents, particularly those powered by large language models and integrated search/payment ecosystems (e.g. ChatGPT, Google Gemini, Klarna AI), do not scroll through pages of visually branded storefronts. They parse structured data (availability, delivery speed, verified reviews, specifications) and optimize for best-fit criteria. In this context, the emotional and narrative components of a brand, long central to B2C marketing strategies, risk being lost in translation unless they are made explicit, quantifiable, and embedded into machine-readable formats.

In B2B, the challenge is even starker. Agentic systems managing procurement or vendor selection will evaluate suppliers on efficiency, compliance, ESG alignment, pricing, and historical reliability—often in milliseconds. Differentiation here can no longer hinge on relationships alone. It must be operationalized into data.

What makes an experience "unique"? In the context of agentic commerce, a unique experience is:

- Difficult to replicate by competitors
- Valuable to both the end user and the interpreting agent
- Rooted in proprietary assets and brand ethos
- Visible in structured form, not just through aesthetics or messaging

Examples include:

- Proprietary recommendation engines that blend behavioural, contextual, and psychographic data
- Immersive product visualization tools, such as AR try-ons or interactive 3D configurations
- Unique fulfilment experiences, such as same-day hyperlocal delivery or carbon-offset logistics
- Data-backed service guarantees, like uptime SLAs in B2B SaaS or AI-assisted customer service routing
- Custom product bundles or subscriptions configured dynamically based on past usage, lifecycle needs, or organizational purchasing behavior

Making differentiation legible to agents

To succeed in an agent-mediated marketplace, unique experiences must be translated into signals that autonomous systems can interpret and act upon. This requires more than just compelling design or compelling copy. It calls for a deliberate restructuring of how businesses surface their differentiators. Product and service features must be explicitly linked to decision-relevant attributes so that agents can evaluate and prioritize them.

For B2C brands, this might involve marking up trust-related elements in structured, machine-readable formats such as verified sustainability credentials, flexible return policies, or customer service standards. It could also mean feeding proprietary insights, such as dynamic sizing recommendations or personal style compatibility, into third-party recommendation engines or shopping assistants. Even loyalty programmes and experiential services like concierge support need to be encoded in a way that agents can detect and factor into their choices.

In the B2B context, the task is no less urgent. Service-level differentiators (e.g., delivery guarantees, bespoke contract terms, or compliance certifications) should be made accessible via procurement APIs or embedded within digital marketplaces. Likewise, the substance of what sets a supplier apart must be reflected in structured formats such as RFP platforms, contract metadata, or digital product passports. The objective is to ensure that differentiation is not just present, but machine-legible and visible to the systems making or influencing commercial decisions at scale.

Change management: Aligning strategy, systems, and culture

The adoption of agentic commerce, while promising in its potential, presents a set of non-trivial challenges for most businesses. These challenges are not solely technological but also organizational, structural, and cultural.

Leaders must contend with the realities of shifting workflows and a hesitation to place trust in autonomous decision-making systems. Teams may be reluctant to delegate high-stakes decisions to systems whose reasoning they cannot easily interpret or audit. This is especially true in regulated industries or customer-facing contexts where transparency is essential.

In addition, technological integration emerges as a critical barrier. Many companies still operate on legacy infrastructure or within tightly coupled systems, which were not designed with real-time intelligence in mind. This creates friction for deploying AI-driven agents, whose autonomy depends on the ability to access modular services via well-defined APIs.

Sources of friction cluster around four key themes:

- Legacy systems that restrict API access and slow integration
- Organizational inertia and anxiety around role displacement
- Low AI fluency across key decision-making units
- Skepticism toward autonomous systems and their decision logic

From friction to enablement

Overcoming these challenges requires a phased and strategic approach. Early implementations should target low-risk, high-value use cases such as content enrichment, campaign personalization, or dynamic pricing where outcomes can be clearly measured and easily reversed if necessary. Agents should operate within clearly defined parameters, with transparent guardrails and human-in-the-loop design to ensure alignment with brand and business goals.

Equally important is organizational readiness. Businesses must invest in cross-functional education, bridging technical and operational teams, and embedding AI literacy into decision-making processes. Feedback mechanisms should be built into agentic systems from the outset, allowing humans to intervene, correct, or re-train agents without obstructing their ability to scale.

Ultimately, the successful adoption of agentic commerce hinges on the synchronization of **technical readiness and cultural enablement**. When approached thoughtfully, these early challenges become opportunities to modernize infrastructure and to reimagine how value is created across the organization.

Agentic readiness: Preparing for what's next

Agentic commerce is no longer a question of **if** but **when** and more importantly, **how**. While the technology is maturing rapidly, its strategic adoption remains uneven. Businesses that move early, with clarity and discipline, will keep pace and define the new rules of engagement. Those that delay, risk becoming irrelevant in a market where autonomy is expected. The path forward is not radical disruption but deliberate transformation, starting with foundational readiness and expanding through iterative deployment. The goal is not to replace human decision-making with machines, but to shift the burden of routine decisions to intelligent systems, **freeing human talent to focus on creativity, strategy, and relationship-building**.

Essential steps to prepare for agentic commerce:

1. **Audit data maturity and eliminate silos**
   Autonomous agents require context. If your data is fragmented, outdated, or inaccessible, agents will act on flawed assumptions. Mapping, cleansing, and consolidating your data landscape is a prerequisite to deploying any meaningful level of intelligence.
2. **Design for transparency and user override**
   Agents should not operate as black boxes. Businesses must design systems that explain their logic in human-readable terms, offer override mechanisms when needed, and build confidence through clarity. Explainability is an ethical safeguard and a strategic advantage.
3. **Start small, prove value, scale iteratively**
   Begin with low-risk, high-impact use cases such as AI-generated product descriptions, intelligent search assistants, or dynamic pricing optimizers. Use these initial deployments to measure value, refine governance, and build internal momentum.
4. **Invest in people. AI fluency is a leadership imperative**
   As agents take over operational tasks, the role of human teams will evolve. Employees must be equipped not only to work alongside autonomous systems, but to guide, evaluate, and improve them. This requires cross-functional training, new mental models, and leadership committed to continuous learning.

Reframing control: Delegating execution, not responsibility

The purpose of agentic commerce is not to replace human decision-making but to reassign it, carefully and selectively, to systems capable of managing operational complexity at scale. It is intent delegation with intelligent oversight. Machines handle execution, humans define the boundaries: strategic, ethical, experiential. Agentic commerce is not a feature to be added, but a capability to be built. Organizations that recognize this and act now will be the ones best equipped to operationalize agentic systems and maintain control as autonomy becomes the norm.

Checklist: Preparing for agentic commerce

Audit data maturity and eliminate silos

- ☐ Design agents for transparency and user override
- ☐ Start with low-risk, high-impact use cases
- ☐ Scale iteratively based on proven value
- ☐ Train teams for AI fluency and cross-functional literacy
- ☐ Ensure explainability and human-in-the-loop oversight
- ☐ Embed agents into composable, API-first infrastructure
- ☐ Treat agentic commerce as a strategic capability, not a tool

Chapter 3

What's not there yet: Limitations, gaps, and critical questions

shopware

Governance and ethical frameworks for agentic commerce

As agentic AI introduces new levels of autonomy into digital ecosystems, it simultaneously raises a fundamental question: **who, or what, is responsible for the decisions being made?** The deployment of commercial agents cannot occur in a vacuum. It must be accompanied by robust governance structures and ethical frameworks that define the boundaries of their behavior, ensure accountability, and preserve trust. Autonomous systems operate at the intersection of data, logic, and intent. While this offers efficiency and scalability, it also shifts decision-making into less visible, more automated layers of the tech stack. In such an environment, **governance is not an afterthought but a design principle.**

Four foundational tenets must guide the responsible deployment of commercial agents:

1. **Guidance**
   Autonomy must have boundaries. For agents to function responsibly in commercial environments, businesses need to define a clear framework for operational authority. What decisions may an agent execute autonomously? Where is human input or review required? Guidance includes not only access rights and data scopes, but also the definition of critical checkpoints—especially where risk, regulation, or brand exposure is involved. For example, an agent that supports merchants in storefront operations might be allowed to propose price changes, but not publish them without review. Draft-and-approve workflows ensure control without stifling automation. Rather than viewing human oversight as a bottleneck, guidance treats it as an integral layer of governance, enabling businesses to scale agentic capabilities with intention.
2. **Explainability**
   Agents must provide transparent reasoning behind their actions. If a system acts on your behalf, do you understand how and why it made that decision? Without explainability, even well-performing agents will struggle to gain user confidence. Logging decision paths, model outputs, and source data is essential—not only for auditability, but also for enabling human oversight. If a customer receives an unusual product recommendation or a dynamic price shift, businesses must be able to trace the logic, not just the output.

3. **Brand alignment**
   While agents may optimize for outcomes (conversion, margin, engagement) they must do so within brand-defined parameters encoded into their logic. Left unchecked, an agent could select a product or message that undercuts years of brand equity in favour of a marginal gain. Tone of voice, visual identity, and customer experience principles are essential strategic assets. An agent operating without these constraints risks flattening differentiation into pure utility.

4. **Interoperability**
   Autonomous agents thrive on access: to data, services, and context. When businesses become overly dependent on proprietary ecosystems or black-box models, they risk fragmenting the user experience and forfeiting control. When each vendor builds their own agent ecosystem with incompatible standards, the result is a fractured landscape. This not only limits flexibility—it recreates the integration headaches composable commerce was designed to eliminate. Open, standards-based architectures reduce this risk by ensuring agents can act seamlessly across systems without creating long-term vendor lock-in.

These principles are not just theoretical. They address risks that are already emerging in early deployments. As MIT researcher and AI ethicist Dr. Sandra Wachter observes, ***"[...] nearly always, all the attention is on what can be done with technology, and no one stops to ask whether it actually should be done."*** (source) Agents are only as ethical as the frameworks in which they operate. If commerce defines the rules, we must ensure those rules have guardrails.

The implication is clear: **Intelligence without structure is not innovation, it's liability**. In this new paradigm, governance becomes a competitive differentiator. Businesses that embed ethical clarity and operational transparency into their agentic systems will reduce risk and build the trust necessary to scale these systems responsibly and sustainably.

The unresolved challenges of agentic commerce

While agentic AI promises a decisive shift in how digital transactions are initiated and executed, its readiness for full-scale commercial deployment is far from assured. Beneath the architectural elegance and strategic promise lies a set of practical, infrastructural, and ethical constraints—many of which remain unsolved. Recognizing these boundaries is not an argument against agentic commerce, but a necessary step in shaping it responsibly.

Energy cost of intelligence

The intelligence that powers agentic systems is not immaterial. Behind each real-time decision, product recommendation, or task automation lies a vast matrix of GPUs and data centers consuming ever-growing amounts of electricity. The AI models underpinning agentic commerce (particularly those capable of learning from feedback, navigating ambiguity, or simulating multi-step outcomes) are computationally intensive.

The industry is reaching an inflection point where the scalability of agentic systems will be constrained not by innovation, but by energy infrastructure. Hyperscalers are exploring nuclear options to support the load. Without breakthroughs in energy-efficient model design, the operational costs of agentic AI may prove unsustainable, particularly for small and mid-market retailers. In this light, the true winners may be those who build leaner, not just smarter, agents.

Reinforcement learning: Essential but imperfect

Reinforcement learning (RL) is often heralded as the enabling logic behind agentic AI. It allows agents to navigate environments, learn from outcomes, and optimize behaviour over time. But while RL marks a departure from static, rule-based systems, it is far from a turnkey solution.

Its limitations are well documented: high training costs, poor transferability across domains, lack of interpretability, and a significant reliance on synthetic data to simulate edge cases. RL is indispensable for end-to-end task execution, but it struggles when context shifts or where outcomes depend on tacit knowledge, ethics, or subjective judgement.

To address this, many AI innovators are now hybridizing RL with supervised learning and retrieval-augmented generation (RAG). But these architectures remain fragile, and most are not yet production-ready for high-stakes commercial contexts.

The data paradox

Agentic commerce is, at its core, data-dependent. And yet, the very thing it requires most (high-quality, domain-specific, real-time data) is the most difficult to guarantee. In fields like healthcare, finance, and ecommerce, legacy data systems are often siloed, incomplete, or governed by inconsistent standards.

Many businesses cite data quality and readiness as their biggest barrier to deploying AI. Without data that is both accessible and trustworthy, even the most advanced agent will falter. In highly regulated sectors, data governance isn't just a technical issue but a compliance risk.

To overcome this, organizations must prioritise foundational data modernization before leaping into agentic workflows. That includes unifying silos, establishing semantic interoperability, and building infrastructure for continuous, real-time updates. Without it, autonomy becomes little more than automation with better branding.

The human factor

Perhaps the most understated limitation of agentic commerce is human readiness. The transition from human to machine-led challenges not only operations, but culture and governance. Many businesses remain hesitant to delegate complex decisions to systems they do not fully understand, especially when those decisions touch on customer relationships, financial strategy, or brand voice.

Agentic commerce may one day handle the full span of a transaction, from discovery to dispute resolution. But for now, it thrives in highly structured, repetitive processes and not in domains requiring empathy, trust, or negotiation. AI for tasks, humans for judgment. This principle may not be permanent, but for the near future, it must remain the organizing logic for responsible adoption.

Chapter 4

Shopware's strategic direction for agentic commerce

shopware

Our vision and roadmap for agentic commerce

In an industry where the pace of innovation often outstrips meaningful implementation, Shopware takes a distinctly pragmatic approach to the future of agentic AI. While the potential of agentic commerce is vast, it's the methodical, practical application of this technology that will determine its success. For Shopware, agentic commerce is not a theoretical aspiration but an evolving roadmap grounded in real-time data and open technology.

At the centre of this strategy is **Shopware AI**, our rapidly expanding suite of AI-driven features designed to assist merchants in everyday tasks. However, Shopware AI is more than a collection of individual tools. It's a foundational capability that powers and connects the entire Shopware ecosystem, across every solution and touchpoint.

Copilot: From automation to autonomy

Shopware Copilot represents our first step into the era of **agentic commerce**, where autonomous software agents move beyond task support to actively shape and optimize workflows. This is not simply a product enhancement. Copilot reflects a broader strategic direction: one where agents operate with initiative, adapting in real time to context, priorities, and customer needs.

Open infrastructure, merchant-centric innovation

Importantly, Shopware's vision does not chase novelty for its own sake. Agentic commerce, in this context, is not a product. It is an **outcome**, made possible by an architecture that is open, extensible, and API-driven. Our approach ensures that merchants are not locked into rigid vendor ecosystems but can compose their own intelligent commerce workflows using the tools that suit their business best. By building on this open foundation, Shopware is also lowering the barrier to innovation. Instead of requiring merchants to invest in expensive proprietary AI stacks, we are embedding intelligence directly into our existing infrastructure, helping businesses scale what matters most: customer relevance and operational efficiency.

Agentic Commerce Alliance

In May 2025, Shopware co-founded the Agentic Commerce Alliance (ACA) to help lay the groundwork for a new era of autonomous ecommerce. The ACA brings together AI-first software companies, researchers, and merchants with a shared mission: to build open, self-optimizing commerce infrastructure. Central to this initiative is the launch of agentic-commerce.org, a platform dedicated to advancing open standards for agent-compatible interfaces and interoperable data exchange. These standards are essential for enabling AI agents to seamlessly access and act on real-time product, inventory, and transactional data across ecosystems. By establishing a common technical foundation, the ACA seeks to reduce operational costs, safeguard data sovereignty, and unlock new revenue models, empowering merchants everywhere to thrive in a more intelligent, decentralized, and scalable commerce landscape.

Looking ahead: The road to practical, impactful agentic AI

The commercial landscape is entering a new phase, defined not by more tools but by smarter orchestration. Shopware's roadmap is aligned with this shift. Adoption will be iterative. Use cases will mature in waves. But with a merchant-first mindset, technical transparency, and a commitment to open innovation, Shopware is laying the groundwork for agentic commerce to move from theory to practice: on real storefronts, solving real problems, in real time.

Conclusion:
Progress with perspective

Agentic commerce is no longer theoretical. Its foundational components—autonomous recommendation systems, real-time pricing engines, and intent-driven user journeys—are already reshaping digital ecosystems. But full orchestration, where agents operate seamlessly across systems and decisions, remains emergent. The promise is real, but so are the limitations.

The path forward demands clear-eyed pragmatism. The challenge is not whether agentic commerce can work, but under what conditions it should operate, what safeguards must govern it, and how it can scale sustainably—technically, ethically, and economically. Energy efficiency, data quality, learning limitations, and human trust are central to making autonomy viable at scale.

Agentic commerce is not a product, nor a feature to be toggled on. It is a structural shift: from interaction to delegation, from user interface to system intelligence. Its transformation will unfold step by step, but its long-term consequences will be generational. Businesses that engage early, with strategic discipline and architectural foresight, will define the next decade of digital commerce.

At Shopware, we're building toward this future. Not as a distant concept, but as a defined and actionable roadmap. If you're exploring what agentic commerce could mean for your organization, across operations, user experience, or business logic, we welcome the opportunity to begin that conversation with you. Start with a personalized demo. Get in touch today. Find our contact details on the next page.

Further reading & resources

Gartner: Future of Sales 2024
MIT Sloan: The Ethics of Autonomous Agents in Commerce
LangChain Documentation
OpenAI: AutoGPT and Beyond
Is Agentic AI Ready To Handle The Way We Do Business?
Forrester's Top 10 Emerging Technologies For 2024: As AI Dominates, Security Becomes Paramount
Agentic AI: 9 promising use cases for business
Professor Sandra Wachter: Addressing the legal and ethical impacts of AI
Open Problems and Fundamental Limitations of Reinforcement Learning from Human Feedback | Montreal AI Ethics Institute
The Agentic Wave is Here—Are You Ready?

Next steps.
New opportunities.

Contact us | Request demo

shopware AG
Ebbinghoff 10
48624 Schöppingen

growth@shopware.com
+49 (0) 2555 92885-0
www.shopware.com

About Shopware

Shopware is Europe's leading open-source ecommerce platform, built to address the technical and operational complexities of modern B2B and B2C organizations. Its composable, API-first architecture and flexible deployment options—SaaS, PaaS, or self-hosted—give businesses full control over infrastructure and the freedom to tailor every layer of the stack. With enterprise-grade tooling, more than 3,100 certified extensions, and an engaged global partner network, Shopware enables rapid adaptation to evolving buyer expectations, complex business models, and regi-on-specific commerce requirements.