CLASSY LLAMA WHITEPAPER

Agentic Commerce for Industrial B2B 2026

What, When, and How to Respond to "Agentic Commerce" for the Mid Market Industrial Merchant

Published May 5th, 2026 by Classy Llama. Written and Edited by Greg Tull, Marketing Director

Key Takeaways

- The Day One shortlist is increasingly compiled by AI before any sales conversation — and 95% of winning vendors are already on it (Forrester, *State of Business Buying 2026*).
- 89% of B2B buyers now use generative AI as a primary research source (Forrester); AI-referred commerce traffic grew 393% year-over-year in Q1 2026 (Adobe Digital Insights).
- Agentic commerce will unfold in four phases for industrial B2B; the accessible ROI in 2026 lives in Phases 1 and 2, not the autonomous-agent endgame.
- Most industrial manufacturers have machine-readable data on only 30–50% of SKUs — invisible to AI shopping agents across the long tail.
- Four converging protocols now define the rails: MCP (Anthropic), ACP (OpenAI/Stripe), UCP (Google/Shopify), AP2 (Google).
- Integration count is the strongest predictor of B2B replatforming complexity and agent readiness (Classy Llama *B2B Commerce Replatforming Benchmarks 2026*, 108 engagements, 11 distinct ERPs).
- Agent readiness is a five-layer stack: catalog data depth, schema/AEO, feed and protocol participation, MCP-ready infrastructure, and governance.

The Day One Shortlist Problem

A procurement director at a $200M agricultural equipment dealer needs replacement hydraulic fittings — 3,000 psi working pressure, SAE 37° JIC, zinc-plated carbon steel. She types the spec into ChatGPT. Within seconds, she has a ranked list of three suppliers that match, complete with lead time estimates and distributor locations.

No rep has been called. No catalog has been opened. The shortlist already exists.

This is the actual change in B2B buying behavior that's happening right now — and it has almost nothing to do with fully autonomous AI agents transacting on your behalf. It has everything to do with whether your product data, your specs, your capabilities, and your brand are legible to a large language model when a buyer uses one as a starting point.

According to Forrester's *State of Business Buying 2026* report, 95% of winning B2B vendors are already on the buyer's shortlist on Day One. That statistic predates widespread AI agent adoption. It's about to get more consequential, not less, as AI systems increasingly compile that shortlist before human conversation begins.

This piece is about what agentic commerce actually is, when industrial B2B will feel its different phases, and — most importantly — what to do about it this quarter. The technology is moving fast. The parts of it that matter to marketing directors at mid-market industrial brands are more knowable than the hype cycle suggests.

What Agentic Commerce Actually Is (and Isn't)

The term "agentic commerce" covers three meaningfully different things, and the marketing noise tends to blur them together. Getting the definitions right changes the action.

Type 1: AI Features Inside Your Storefront

Semantic search, conversational product finders, AI-generated descriptions, smart recommendations — these are tools embedded in your own commerce experience to help buyers navigate it. They're useful, increasingly table-stakes, and not agentic in the meaningful sense. An AI chatbot on your site that answers questions about product specs is impressive. It is not an agent. It cannot act on your behalf, has no persistent memory across sessions, and takes no autonomous actions outside of your storefront UI. The buyer is still in the driver's seat.

Type 2: AI-Mediated Discovery

This is where industrial B2B should be paying close attention *right now*. According to Forrester's *State of Business Buying 2026*, 89% of B2B buyers use generative AI as a top source of self-guided information. Adobe Digital Insights' Q1 2026 report tracked a 393% year-over-year increase in AI-referred traffic to US retail and commerce sites. When a procurement professional asks an LLM — ChatGPT, Perplexity, Google Gemini, Microsoft Copilot — which suppliers carry NSF-certified stainless fittings in 1/2″ NPT, the model answers from its training data and from structured sources it can read. Whether your products appear in that answer is a data and discoverability question, not a sales question. It is the new SEO. It is happening now.

Type 3: Agentic Commerce Proper

True agentic commerce involves AI agents acting with delegated authority: researching, comparing, negotiating, purchasing, and fulfilling — autonomously, on behalf of a human principal who has set parameters and approved the scope. The agent doesn't just find the hydraulic fitting. It compares it against three approved vendors from the procurement system, checks the contract price, confirms inventory at the nearest distribution center, raises a purchase order, routes it for approval, and confirms delivery. The human set the rules; the agent executed them.

This is the version that commands the largest predictions. Gartner forecasts that by 2028, 90% of B2B buying could be AI agent-intermediated, representing more than $15 trillion in agent-mediated spend. McKinsey QuantumBlack projects agentic commerce orchestrating $3–$5 trillion in global retail spend by 2030. These are not numbers to dismiss — but they live primarily in Phase 3 and Phase 4 of a trajectory that industrial B2B is still entering.

The Protocol Stack: The Rails Being Built Right Now

Four protocols, developed in 2025–2026, define how agents will interact with commerce systems. You don't need to understand their technical specifications to understand why they matter.

MCP (Model Context Protocol), developed by Anthropic, is the foundational plumbing layer: a standardized way for AI models to read from any tool, data source, or API. Think of it as the language agents use to talk to software. Without MCP, every agent-to-system connection requires custom integration work.

ACP (Agentic Commerce Protocol), co-developed by OpenAI and Stripe (with Meta joining as a supporter), is an Apache 2.0 open standard that defines product feed schema and checkout API endpoints. It powers ChatGPT's product search and ordering capabilities. The spec is versioned and public at agenticcommerce.dev.

UCP (Universal Commerce Protocol), co-developed by Google and Shopify with 20+ co-developers and endorsers — including Walmart, Target, The Home Depot, Wayfair, Visa, Mastercard, Amex, Adyen, and Etsy — is designed to be the broadest interoperability layer. It is explicitly compatible with ACP, MCP, and A2A (the Agent2Agent protocol for agent-to-agent negotiation).

AP2 (Agent Payments Protocol), developed by Google, is a payment-agnostic framework for settling transactions between agents. Mastercard, PayPal, American Express, Coinbase, Adobe, Adyen, and Airwallex all support it.

The significance isn't in any individual protocol. It's that the major players in commerce — platform providers, payment networks, search companies, AI labs — are converging on shared standards. This is how internet protocols work when an ecosystem has reached the "let's stop doing this separately" inflection point.

One important anchor point for merchants: even in agent-mediated transactions, the merchant of record remains you. The customer relationship, the brand trust, the return policy — those stay with the seller. OpenAI's rollback of full in-chat checkout completion in March 2026 (most purchases now redirect to the merchant storefront via in-app browser) reinforces this reality. Your owned checkout experience still matters even as discovery becomes agent-mediated.

A Candid Reality Check

Consumer AI shopping demand is real, but the infrastructure hasn't caught up. OST Agency ran a controlled test with ten mid-market ecommerce stores — a mix of Shopify, WooCommerce, and custom platforms in the $2M–$50M revenue range, across apparel, home goods, beauty, and outdoor gear — asking a Comet-based shopping agent to complete a purchase on each. Only three completed end-to-end. Four were blocked at the product page by modal popups, newsletter signup overlays, or cookie banners the agent couldn't reliably dismiss; the remaining three failed on JavaScript-dependent pricing that wasn't exposed to feeds, CAPTCHA gates, and inventory mismatches between the product feed and the live cart. Separately, academic research by Kaiser and Schulze — analyzing 973 ecommerce sites with $20 billion in combined revenue, comparing approximately 50,000 ChatGPT referral transactions against 164 million transactions from traditional channels — found that affiliate links convert 86% more often than ChatGPT referrals, and organic search outperforms ChatGPT by roughly 13%. The gap isn't about buyer intent; it's about merchant infrastructure not being built for agents.

B2B has additional structural complexity that makes the gap wider. According to Digital Commerce 360's analysis: "For B2B, the near-term opportunity is AI-assisted workflows rather than fully autonomous purchasing." That framing is accurate, and it's useful: the near-term opportunity is large enough to act on now, even if the autonomous-agent endgame is still several years away.

"For B2B, the near-term opportunity is AI-assisted workflows rather than fully autonomous purchasing." — Digital Commerce 360

When Industrial B2B Will Feel This

Industrial B2B won't experience agentic commerce as a single transition. It will unfold in phases, and different phases call for different responses. Rejecting the binary "is it here or not?" framing lets you act on what's actually useful today.

Phase 1 — AI-Mediated Discovery (Now, 2026)

This phase is not coming. It's here.

Forrester's *State of Business Buying 2026* documents what the B2B buying process now looks like: an average of 13 internal stakeholders and 9 external influencers involved in each purchase decision, with 89% of buyers using generative AI as a primary source of self-guided research. Procurement is a decision-maker in 53% of buying cycles and is engaged from the start. AI at Wharton's "Growing Up: Navigating Gen AI's Early Years" study found that 94% of procurement executives use generative AI at least weekly — the highest weekly adoption rate of any business function surveyed.

For industrial manufacturers and distributors, this creates a structural problem: most of what's in your catalog isn't legible to an AI. Industry research suggests that most manufacturers publish structured content on only 30–50% of their SKUs — the long tail is missing attributes, lacks machine-readable specs, and is invisible to the models building buyer shortlists. A chemical supplier with 8,000 SKUs that has detailed, structured data for its top 200 products but unformatted PDFs for the rest is invisible to AI for 97% of its catalog.

The Forrester Day One statistic sharpens the stakes: if 95% of winning vendors are on the buyer's initial shortlist, and that shortlist is increasingly being compiled by an AI before any human outreach begins, then the discovery problem is the top-of-funnel problem. It doesn't matter how good your sales team is if you're not in consideration.

Phase 2 — AI-Assisted Ordering on Owned Channels (2026–2027)

The most actionable opportunity for mid-market industrial brands over the next 18 months isn't in third-party agent platforms. It's in deploying AI assistance on your own B2B portal to make ordering faster and more accurate for your existing customers.

Customer-facing conversational agents on a B2B portal — one that understands contract pricing, negotiated terms, approved SKUs, and a buyer's order history — can materially reduce the friction of reorder cycles, handle complex spec-matching queries, and surface substitutions when items are backordered. This is where the ROI is

clearest and the required infrastructure is closest to what many mid-market industrials already have.

The prerequisite is ERP integration that actually works. In our replatform engagements with mid-market industrial brands, ERP integration is consistently the longest lead-time item and the most common source of scope creep. Across the 108 B2B commerce projects in Classy Llama's *B2B Commerce Replatforming Benchmarks 2026* report, projects with five or more integrations took 12+ months without exception. The data also surfaces a distribution of 11 distinct ERP systems across named projects alone — NetSuite, Infor SyteLine, Infor M3, Microsoft Dynamics 365, Microsoft Business Central, Epicor, Sage, Acumatica, Agility, Fishbowl, and several unnamed or custom-built systems. There is no single dominant ERP in B2B mid-market, which means there is no off-the-shelf AI ordering layer that works with your system without real integration work underneath it.

The B2B portal opportunity is real and measurable. Russell Hendrix — Canada's largest food service equipment supplier, operating for 80 years — documented the impact in a published Shopify case study: after migrating from a legacy platform to Shopify B2B, they saw a 24% increase in revenue, a 43% increase in B2B online order volume, and a 9% increase in average order value in the 12 months post-launch. Order processing speed improved 5x. Their platform now manages 300+ customer groups across more than 10,000 products. That baseline — faster ordering, higher attach rates, more volume through the digital channel — is the foundation on which AI-assisted reorder and conversational ordering layers will be built.

At the enterprise end of the industrial spectrum, Carrier — the global HVAC, refrigeration, and building controls manufacturer operating across 180 countries — documented in a Shopify case study that their previous monolith ecommerce platform required 9–12 months and up to $2 million per site to launch a new regional commerce experience. On Shopify, those figures dropped to 30 days and approximately $100,000 per site — a 90%+ reduction in both time and cost. Their OneCommerce accelerator now covers approximately 80% of shared B2B commerce requirements across markets. The implication for industrial brands considering agent-readiness investment: infrastructure that can be deployed and updated rapidly is infrastructure that can adapt to agent protocols as they mature.

For the search rescue pattern that typically precedes AI-readiness work, a mid-market sporting optics manufacturer in our portfolio illustrates a challenge common in complex B2B catalogs: after deploying Algolia, the team discovered that semantic search was returning irrelevant results not because the search technology was wrong but because the underlying product attribute data was too sparse and inconsistently structured for the engine to work with. Rebuilding the data layer — normalizing attribute values, completing the long-tail SKU specs, establishing controlled vocabularies for product types — was what fixed it. The search UI was the symptom; the catalog was the problem. That same data quality gap is the reason AI agents misfire on industrial catalogs: the agent's output is only as good as the structured data it reads.

Phase 3 — Procurement-Side Agents Negotiating With Seller-Side Counterparts (2027–2029)

Forrester predicted in early 2026 that by year-end, 1 in 5 B2B sellers would face AI-powered buyer agents in quote negotiations. The procurement platforms — JAGGAER, Coupa, SAP Ariba, Zip — are actively building this capability. The 2025 ProcureCon CPO Report (produced by Icertis with ProcureCon Insights) found that 66% of CPOs list AI adoption as a top priority, and 90% are already evaluating or using AI agents.

When this arrives, the implication for sellers isn't just about having good prices. It's about having machine-readable pricing logic, negotiation parameters, and credit/terms structures that agent-side systems can interact with. A seller whose pricing lives in a spreadsheet that a rep adjusts manually cannot respond to an AI-negotiated RFQ at the speed the buyer's system expects.

The quality problem is real. Forrester found that 28% of procurement respondents felt less confident in a decision because of inaccurate AI output — a meaningfully higher rate than among all buyers (19%). Twenty-two percent wasted time because of poor AI information. The practical takeaway: the AI agents coming for B2B negotiations are not infallible, and sellers should build their systems to handle AI-initiated negotiations with appropriate guardrails, not assume the other side's agent is correct.

Phase 4 — Autonomous Agent-to-Agent Transactions for Catalog Purchases (2028+)

The $15 trillion Gartner projection lives here. MRO reorder is the realistic near-term use case: consumables with clear specs, known suppliers, established contract terms, and no configuration complexity. When all of those conditions hold, the case for keeping a human in the reorder loop weakens.

Most strategic, configured, or contract-priced industrial purchases will keep humans in the loop for the foreseeable future. A $2M piece of agricultural equipment, a custom formulation order, a first-time purchase from a new supplier — these involve judgment that procurement professionals are not delegating to agents in 2028.

The honest take: The dollars at stake in Phase 1 (being on the shortlist) and Phase 2 (assisted ordering on your own portal) are accessible now. They require data work, integration investment, and schema implementation — not a leap to autonomous agent infrastructure. Most of this article's "How" is about Phases 1 and 2.

How to Get Ready: A Five-Layer Framework

Industrial B2B's structural complexity — ERP-tied pricing, contract terms, customer-specific SKUs, real-time inventory across multiple locations, spec-heavy catalogs, distributor channel overlaps — means agent readiness is not a toggle to flip. It's a stack to build. Each layer below includes what it is, what we typically find in mid-market industrial brands when we start replatform discovery, and what to do about it.

Layer 1 — Catalog Data Depth

Bluestone PIM's published benchmark defines 30+ structured fields across six categories that AI shopping agents evaluate before recommending a product. For D2C, those fields are dominated by visual attributes, size, color, and reviews. For industrial B2B, the relevant fields look different:

- Technical specifications: Dimensional tolerances, pressure ratings, temperature ranges, material composition (grade, alloy, ASTM/ASME standard)
- Certifications and compliance: UL listing, ISO certification, RoHS compliance, NSF approval, FDA contact-safe status
- Application attributes: Chemical compatibility (specific fluid, concentration, temperature), industry code or classification (SIC/UNSPSC), fitment data (OEM cross-reference, year-make-model for replacement parts)
- Logistics attributes: Freight class, package dimensions, hazmat classification, country of origin
- Channel attributes: GTIN/EAN, manufacturer part number (MPN), distributor SKU cross-reference, lead time by fulfillment location

The gap we consistently find in replatform discovery: manufacturers have these attributes in engineering systems, PLM software, or product datasheets — but they were never structured and loaded into the commerce catalog. Products with "Weatherproof: Yes" in a freeform field are invisible to agents looking for IP_rating: IP67. Products with applications described in marketing prose ("great for demanding industrial environments") cannot be matched to a spec-based search.

The 30–50% SKU coverage gap in industrial catalogs isn't a data entry problem — it's a structural integration gap between engineering systems and commerce systems.

Machine-readable attributes beat marketing copy in every AI evaluation. "Waterproof rating: 15,000mm" surfaces in structured queries. "Exceptionally waterproof performance" does not. For industrial: "Burst pressure: 6,500 psi, carbon steel, NPT 3/4-inch" is what agents read. "Built for demanding applications" is not data.

What to do: Audit your top 500 SKUs for structured attribute count. Below 15 meaningful technical attributes per product is a gap you can prioritize closing. Identify which attributes live in ERP, PLM, or product engineering systems and are not flowing to commerce — that's the integration work that unlocks the long tail.

Layer 2 — Schema and Answer Engine Optimization (AEO)

Structured data markup tells AI systems what your page is about without requiring them to infer it from prose. The baseline schema types for commerce — Product, Offer, AggregateRating — are table-stakes. The ones that specifically matter for AI agent eligibility and citation in AI-generated answers include:

- Product with MPN and GTIN: Allows agents to cross-reference your product against known part numbers in procurement systems
- Offer with real-time pricing and availability feeds: Agents evaluating sourcing options need to know price and stock status — static markup is better than nothing; live feed integration is better still
- FAQPage: Answer engine optimization (AEO) depends on question-answer structure. Procurement professionals ask specific questions: "Is [product] compatible with hydraulic oil ISO 46?" If your site answers that question in a structured FAQ, AI systems can surface your answer in generated responses
- MerchantReturnPolicy and OfferShippingDetails: Required schema types for eligibility in Google's merchant surfaces; increasingly evaluated by agents assessing purchase feasibility

Beyond schema, AEO is about building content that answers the specific questions industrial buyers ask before they've talked to anyone. Application guides, compatibility matrices, certification documentation, installation specs — structured, linked, and authored with enough topical authority that AI systems treat your site as a primary source rather than a secondary one. The same structural investment that improves your AEO position is what makes your content citable in AI-generated buying guides.

When Classy Llama runs AEO audits against mid-market industrial sites, the most common finding isn't missing schema (though that's common too) — it's that schema exists for core catalog pages but is absent or malformed on application-specific content, technical documentation, and long-tail product variants. A product page for a 3/8" fitting might have schema; the page answering "which fitting is rated for 3,000 psi with 316 stainless steel" often doesn't exist at all.

Layer 3 — Feed and Protocol Participation

With the protocol stack now established, the practical question for industrial brands is which surfaces to enable, in what order, and at what scope.

A decision framework:

If you're on Shopify Plus: Agentic Storefronts were activated by default for eligible stores in March 2026. From a single admin toggle, your catalog syndicates to ChatGPT, Microsoft Copilot, Google AI Mode, and Gemini, with full attribution per channel. This is the lowest-friction path to protocol participation available today. Enable discovery first; evaluate transactional capability against your checkout complexity and customer-specific pricing requirements before enabling agent checkout.

If you're on Adobe Commerce: Adobe shipped its Commerce MCP server at Summit 2026, covering catalog, cart, pricing, inventory, promotions, checkout, OMS, and post-purchase. Adobe has committed to UCP, ACP, and AP2 support. The Brand Concierge capability provides an owned conversational commerce experience for companies that want to control the agent interaction rather than surrender it to third-party platforms. Adobe powers approximately 250,000 live commerce sites globally, and the MCP server means existing Adobe Commerce builds are closer to agent-ready infrastructure than they may appear.

If you're on BigCommerce: BigCommerce's API-first architecture supports the same catalog and checkout outcomes as the protocol-native platforms. It was not a co-developer on UCP or ACP, but implementation through the API surface is achievable. The gap is operational: feeds and protocol endpoints require more configuration work than the platform-native paths above.

For all platforms: Enable discovery-ready feeds before transactional. The OpenAI Instant Checkout rollback in March 2026 — which redirected most in-chat purchases back to merchant storefronts — is a relevant data point: agent discovery is more mature and more reliable than agent checkout. Phase the rollout to match infrastructure maturity.

Note on B2B-specific complexity: For industrial brands with customer-specific pricing, contract-gated catalogs, and credit-term-based checkout flows, transactional agent

enablement requires careful design. An agent that browses your catalog and surfaces products to a buyer is useful and low-risk. An agent that attempts to complete a purchase at list price for a customer whose contract price is 40% lower is a problem. Solve the pricing exposure question before enabling transactional flows.

Layer 4 — MCP-Ready Commerce Stack

This is where industrial B2B diverges most sharply from general commerce agent readiness — and where deep ERP integration expertise matters more than any protocol.

Agent-readiness for B2B isn't primarily about having a clean product feed. It's about whether your commerce stack can securely expose the right data to authorized agents:

- Contract pricing and customer-specific catalogs: A distributor with 400 active accounts, each on different negotiated terms, cannot expose a single catalog to agents without account-level authentication. Any agent integration that surfaces contract pricing must do so only to authenticated buyer-side systems with appropriate permissioning.
- Real-time inventory: Agent-initiated purchasing fails when the feed says "in stock" but the cart says "3-week lead time." Real-time inventory feed integration — across multiple warehouse locations if applicable — is prerequisite to transactional agent flows. In our B2B builds, real-time multi-location inventory is a standard requirement in Tier 2+ builds; ErieTec and C&B Equipment are examples in our portfolio where this was a core integration requirement.
- Approval workflows and credit limits: B2B purchasing is not one buyer with a credit card. It involves authorization hierarchies, spend limits by buyer level, and credit terms that determine whether a purchase can happen at all. Agents that bypass these workflows create compliance problems. Agents that work within them require that the workflows be machine-accessible, not just UI-accessible.
- ERP as the system of record: Every one of the above — pricing, inventory, credit, catalog entitlement — ultimately lives in the ERP. If the ERP-to-commerce integration is one-directional (ERP pushes data on a nightly sync), the commerce layer cannot support real-time agent interactions. Bidirectional, near-real-time ERP sync is the foundation any meaningful Phase 2 or Phase 3 capability requires.

Across our portfolio of 108 more recent commerce engagements, the diversity of ERP systems in play — 11 distinct named ERPs, Celigo appearing as integration middleware in 5+ projects alone — underscores that this is custom engineering work in almost every case. NetSuite, Acumatica, Business Central, Epicor, and Infor SyteLine each expose data differently, require different authentication patterns, and have different capabilities for near-real-time sync. Pure-play Shopify builds without ERP fluency hit a ceiling here: the storefront can be beautiful and the catalog can be schema-perfect, but without a production-grade ERP integration, the data that agents need most — live pricing, accurate inventory, contract entitlements — isn't reliably there.

The Classy Llama *B2B Commerce Replatforming Benchmarks 2026* report documents this pattern directly: integration count is the single strongest predictor of project complexity and timeline. It is also the single strongest predictor of agent readiness. This isn't a coincidence.

Layer 5 — Governance and Guardrails

The capabilities above create decisions that marketing directors cannot make alone. Procurement's involvement — and IT's, and legal's — is not optional once you move beyond feed-based discovery into agent-assisted ordering.

Delegated authority frameworks: Before any agent can transact on behalf of a buyer, the buyer (or their organization) must define: What can the agent buy? Up to what spend limit? From which approved vendors? Under which contract terms? These frameworks are not technology problems — they are procurement policy problems that need to be solved before the technology can implement them.

Audit trails and agent identity: A purchase made by an AI agent on behalf of a human principal creates a paper trail question. Whose approval is attached to the purchase order? How do you distinguish agent-initiated orders from buyer-initiated orders in your OMS? This matters for chargebacks, disputes, fraud investigation, and financial reconciliation.

Fraud and error exposure: AI agents acting with purchasing authority are a new attack surface. Prompt injection — where a malicious actor embeds instructions in product catalog data or checkout flows to manipulate agent behavior — is a documented concern. The Center for Data Innovation's 2026 analysis identified significant gaps in how existing Regulation E consumer protection frameworks apply to agent-initiated transactions. The NIST AI Agent Standards Initiative published preliminary guidance in early 2026; final standards are not yet in place. The practical guidance: don't enable agent purchasing flows without explicit authorization scoping, spending limits, and anomaly detection on order patterns.

Questions to ask your platform vendor before enabling agent checkout:

- Does your agent authentication support customer-specific contract pricing?
- Can you restrict agent-initiated orders to pre-approved SKU lists per account?
- Do agent-initiated orders route through your standard approval workflow?
- What audit logging exists for agent-initiated checkout events?

What This Means for Procurement (and What Marketing Directors Should Know About It)

The framing so far has been from the seller's side — how to show up in AI-mediated discovery, how to serve AI-assisted ordering on your portal. The procurement side of the same transaction is moving at least as fast.

AI at Wharton found that 94% of procurement executives use generative AI at least weekly — the highest weekly adoption of any business function. The 2025 ProcureCon CPO Report found 80% of CPOs name AI a top investment priority. At the same time, Forrester's *State of Business Buying 2026* found that 28% of procurement respondents had felt less confident in a decision because of inaccurate AI output — a higher rate than among all buyers (19%). 22% wasted time because of poor information from AI tools. The gap between adoption and trust is real.

Procurement is becoming the buyer of agentic procurement platforms — JAGGAER, Coupa, SAP Ariba, Zip — at the same time that marketing is becoming the seller-side infrastructure owner responsible for agent-facing data and checkout. These two teams, which rarely talk about technology together, are about to be the two halves of the same transaction.

The marketing director's practical interest here: get into the conversation with procurement before your platform vendor does. The questions that matter — What pricing data needs to be machine-readable? Which SKUs are eligible for agent-initiated reorder vs. requiring human approval? What negotiation parameters can a seller-side counter-agent work within? — are cross-functional decisions, not marketing decisions or procurement decisions alone.

A note on email: industry benchmarks indicate that 50–70% of B2B order and quote volume in manufacturing and distribution still flows through email. Before your company is ready for agent-to-agent transactions, the unsexy work of getting that email-based workflow into a structured system — EDI, API, or at minimum a structured portal — is prerequisite. The same data discipline that makes email orders processable makes them agent-processable.

This is a cross-functional moment. The brands that align marketing and procurement on agent-facing infrastructure in 2026 will have a structural advantage in 2028. The ones that wait for procurement to lead, or for marketing to solve it alone, will find themselves rebuilding from a worse starting position.

What to Do This Quarter

Ordered by impact and feasibility for a $30M–$1B industrial brand:

1. Run an AI visibility audit. Ask ChatGPT, Perplexity, and Google Gemini to recommend products in your category. Use actual buyer queries — the kind your sales team hears. Note where you appear, where you don't, what they say about you, and who they're citing instead. This is a 30-minute exercise that surfaces more signal than a quarter of internal debate. Screenshot the

responses and share them with your leadership team; the visual tends to accelerate urgency.

2. Audit catalog completeness on your top 20% of SKUs. Count structured, machine-readable attributes per product. Not freeform description words — actual attribute fields with controlled values. Below 15 is a problem for AI legibility. Below 8 is a problem for any search, AI or otherwise. Identify which attributes your buyers actually use to specify and which live in your ERP or engineering system but haven't been loaded to commerce.
3. Implement or validate Product/Offer/AggregateRating schema across your catalog. Add FAQPage schema to application-specific content, compatibility guides, and technical documentation. Validate MPN and GTIN consistency — these are the identifiers procurement systems use to cross-reference against known parts. OfferShippingDetails and MerchantReturnPolicy schema rounds out the eligibility baseline.
4. Enable agent-friendly discovery surfaces. If you're on Shopify, enable the Agentic Storefront for discovery (evaluate transactional separately). If you're on Adobe Commerce, the MCP server is available — work with your integration team to scope feed quality and pricing exposure before enabling. If on BigCommerce, scope the API configuration work to participate in ACP and UCP. Discovery-ready first, transactional later.
5. Get your ERP integration story straight. Real-time inventory, contract pricing, customer-specific catalog entitlements — these are non-negotiable for any meaningful Phase 2 or beyond. If your ERP integration runs on a nightly sync or requires manual intervention to reflect pricing changes, that's the foundational work that unlocks everything else. If your ERP integration is brittle, this is where you start — not because of agentic commerce specifically, but because every valuable digital channel capability depends on the same foundation.

Closing: The Shortlist Is Already Being Built

The shift in B2B buying isn't about machines replacing procurement judgment. It's about machines getting to the buying decision before humans do. The shortlist that procurement compiles on Day One increasingly reflects what an AI found before any supplier was called.

The brands that show up on that AI-built shortlist in 2026 — because their specs are structured, their data is machine-readable, and their catalog is legible to the models buyers rely on — are the brands still in consideration in 2028. The brands that wait for full autonomous agent transactions before acting on catalog data quality are missing the phase where the ROI is highest and the work is most tractable.

If you're evaluating the readiness of your commerce infrastructure for this shift, Classy Llama's *B2B Commerce Replatforming Benchmarks 2026* report offers first-party data on what production-grade B2B commerce builds actually require, what they cost, and what outcomes they produce. The *Industrial B2B Replatforming Playbook* translates that data into an actionable engagement framework. Both are available at classyllama.com.

The platform decisions, the ERP integration work, and the catalog enrichment that agent readiness requires are the same investments that improve your portal experience, your search performance, and your sales team's productivity today. You don't have to bet on the agentic endgame to act on the foundation. The foundation just became more urgent.

Sources

Forrester Research *State of Business Buying 2026* (January 2026) — Day One shortlist, genAI adoption among B2B buyers, procurement involvement, AI accuracy concerns

Gartner Top Strategic Predictions for 2026 and Beyond (October 2025) — B2B agent-mediated spend forecast ($15T+ by 2028)

McKinsey QuantumBlack *The Agentic Commerce Opportunity* (October 17, 2025) — Agentic commerce orchestrating $3–$5T global retail spend by 2030

Adobe Digital Insights Q1 2026 AI Traffic Report — 393% YoY increase in AI-referred traffic to US retail sites

AI at Wharton / GBK Collective *Growing Up: Navigating Gen AI's Early Years* (October 2024) — 94% of procurement executives use generative AI at least weekly

ProcureCon / Icertis 2025 ProcureCon CPO Report — 66% of CPOs list AI adoption as a top priority, and 90% are already evaluating or using AI agents

Digital Commerce 360 "Forrester: B2B buying groups expand as they question AI" (January 2026) — "For B2B, the near-term opportunity is AI-assisted workflows rather than fully autonomous purchasing"

Kaiser & Schulze (SSRN) *ChatGPT Referrals to E-Commerce Websites* — 973 sites, $20B combined revenue; affiliate links convert 86% more often than ChatGPT referrals; organic search outperforms ChatGPT by ~13%

OST Agency "We Tested 10 Ecommerce Stores With a Shopping Agent. Here's What Broke" (2026) — 3 of 10 mid-market stores complete agent-initiated purchase end-to-end; failure mode taxonomy

Bluestone PIM "Get Recommended by AI Shopping Agents" — AI shopping agents evaluate 30+ structured fields across six categories before recommending a product

Center for Data Innovation "Agentic Commerce is Coming, but Regulation Meant for Humans Will Slow It Down" (March 2026) — Regulation E gap analysis on agent-initiated transactions

NIST AI Agent Standards Initiative (CAISI, February 2026) — preliminary guidance and standards development

Anthropic Model Context Protocol (MCP) documentation

OpenAI / Stripe Agentic Commerce Protocol (ACP)

Google / Shopify Universal Commerce Protocol (UCP)

Google Agent Payments Protocol (AP2)

Shopify Agentic Storefronts announcement (March 2026); UCP co-developer

Adobe Commerce Summit 2026 announcements — MCP server, UCP/ACP/AP2 commitments, Brand Concierge

Russell Hendrix / Shopify Shopify case study — +24% revenue, +43% B2B online order volume, +9% AOV, 5x faster order processing (12 months post-launch)

Carrier / Shopify Shopify case study — 90% faster launch (30 days vs. 9–12 months), 90% cost reduction ($100K vs. $2M per site), ~80% of shared B2B requirements covered by OneCommerce accelerator

Classy Llama *B2B Commerce Replatforming Benchmarks 2026* — first-party data from 108 commerce engagements (2018–2026), 56 confirmed budget actuals; ERP integration patterns, performance outcomes