

AGENTIC COMMERCE · B2B

# Agent-Ready Without the Rebuild

Agentic commerce will be won on catalog data. For B2B, the fix is a data layer that makes the catalog you already have legible to agents. Not a platform migration.

## AUDIENCE

Manufacturers, distributors, wholesalers

## READING TIME

About 9 minutes

## SERIES

The March of Commerce

01

## The short version

**The data argument is right.  
The migration is the wrong  
answer for B2B.**

Every webinar on agentic commerce opens with the same premise. An AI agent can only sell as well as the data underneath it lets it sell. That is correct. Agents read data, not marketing. If the product information is incomplete, inconsistent, or trapped in a format a machine cannot parse, the agent gives a wrong answer or no answer at all.

Where the conversation goes wrong is the prescription. The standard fix, borrowed from retail, is to rip out the old system, clean the data, migrate to a modern platform, and rebuild before agents touch anything. For a B2B manufacturer or distributor still running on PDFs, phone, and email, that is a six to eighteen month project with real cost and real risk. It is the exact project most B2B companies have avoided for a decade.

B2B's catalog problem is real, and in important ways it is worse than retail's. But the fix is not a migration. It is a layer that reads the catalog you already have, structures it, and exposes it to agents and buyers. This paper explains the difference and what to do about it now.

# 67%

of B2B buyers now prefer a rep-free path. 70% want a fully self-service experience.

GARTNER B2B BUYER SURVEYS, 2026

# 75%

would switch suppliers for a better buying experience. 91% among US buyers.

SANA COMMERCE B2B BUYER REPORT, 2025

# 41%

of manufacturing buyers cannot locate the products they need, even on existing web stores.

SANA COMMERCE B2B BUYER REPORT, 2025



02

## The premise everyone agrees on

### An agent is only as good as the data it can read

This part is not in dispute. An agent does not interpret a glossy catalog or read intent the way a person does. It calls data. When a buyer asks whether a part fits their machine, the agent needs the compatibility rule. When they ask for a price on 500 units, it needs the volume break. When they want to order, it needs the minimum quantity and the lead time.

If those facts are missing or unstructured, the result is predictable. Failed pilots, automation that gives confident wrong answers, and rework. Garbage in, failed agent out. So the data foundation matters. The only real question is what fixing the data actually requires.



### 03 Why B2B is harder than retail


## Retail polishes a feed. B2B has to recover knowledge from documents and people.

In retail, the product data usually already lives in a structured feed. A commerce platform holds a title, a price, a few variants, some images. It is often messy. But it starts machine-readable. The work is cleanup.

B2B is a different problem. The product knowledge that actually closes a sale lives in places no agent can read:

- ◆ A PDF catalog where the spec table is an image, not data
- ◆ A stack of product documents that hold the real answers: datasheets, fitment and installation guides, safety sheets, full manuals
- ◆ A spreadsheet a sales engineer keeps on their own desktop
- ◆ Pricing logic that changes by customer, volume, and contract
- ◆ Minimum order quantities, units of measure, and lead times that vary by SKU
- ◆ Compatibility and configuration rules that live in a rep's head

This is the part retail-built advice skips. A B2B catalog is not a clean feed waiting to be polished. It is institutional knowledge spread across documents and people. That is a harder data problem, and it is why a generic "migrate your store" plan does not fit the business it is being sold to.



## 04 Why the standard prescription fails

### For a company on PDFs, “modernize first” means the project they already refused

The default roadmap is some version of this: audit, clean, map, migrate to a modern platform, then deploy agents on top. For a retailer already on a platform, that is a data hygiene exercise. For a B2B company still quoting by phone and email, the same roadmap quietly contains a full eCommerce build before they reach the part they wanted.

That is a long project with migration risk, a heavy bill, and a long wait before any return. It is the reason most B2B companies have stayed where they are. Roughly 60 to 80 percent still sit in the early stages of commerce maturity, working from catalogs and forms. Telling them the price of agentic commerce is the migration they have avoided for ten years guarantees another year of nothing.

#### THE REFRAME

Agent-readiness is a property of the **data**, not the platform it happens to sit on. You can make a catalog legible without moving it.



05

## What agent-legible actually requires

### Strip away the platform talk and the requirement is specific

An agent needs to answer a buyer the way a strong sales rep would. To do that, the catalog data has to expose a short, concrete list:

- ◆ Every product with its real attributes, not a picture of a spec sheet
- ◆ Pricing it can apply, including volume breaks and customer-specific terms
- ◆ Minimum order quantities and units of measure
- ◆ Lead times and availability
- ◆ Compatibility and configuration rules
- ◆ The ability to turn all of it into a quote or an order

Now notice what is not on that list. A new storefront. A platform migration. A re-platformed ERP. None of those are requirements for agent-readiness. They are requirements someone added because that is what they sell.

#### **Structured fields answer "what." Documents answer "will it," "is it," and "how."**

The list above is the structured layer. It tells an agent what a product is. But the questions that actually stall a B2B sale live in documents, not fields. Will this part fit my model. Is this material safe to ship alongside that one. What is the torque spec. How do I install it.

Those answers sit in datasheets and spec sheets, fitment and installation guides, safety data sheets, and full product manuals. A real sales engineer pulls from all of them without thinking about it. An agent has to read them too. So agent-readiness is not only structured attributes. It is the extended product knowledge a buyer needs before they commit, made legible alongside the catalog.



## 06 A layer, not a migration

### Make the catalog answerable instead of moving it

This is the move the retail framing misses. You do not have to relocate the catalog to make it legible. You add a layer that reads what already exists, structures it, and exposes it to agents and buyers.

Ingest the catalog and the documents around it. Pull the attributes, the pricing logic, the order rules, and the answers buried in datasheets, fitment guides, safety sheets, and manuals. Hold them in a structured, machine-readable form. Serve them through a buying experience a person can use, and through endpoints an agent can call. The source catalog does not move. The website does not get rebuilt. The data becomes legible in days, not quarters.

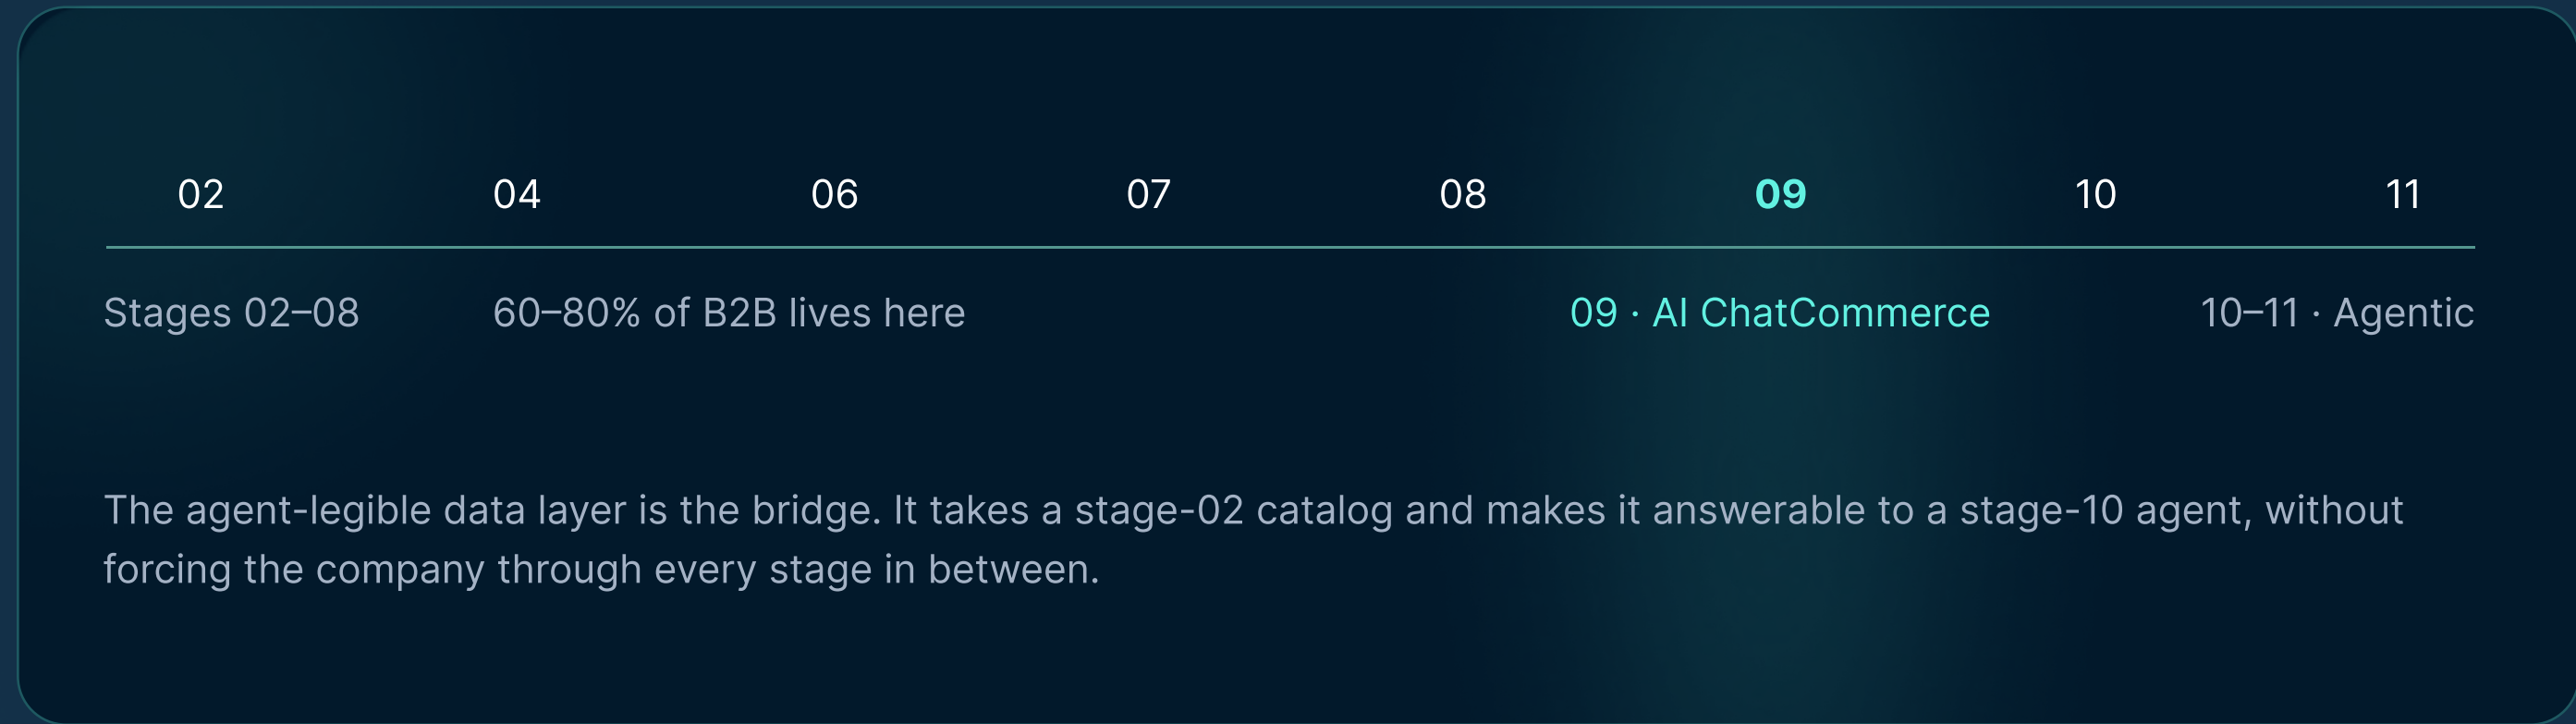
**The distinction that matters.** Modernizing the system and making the system answerable are two different projects. B2B does not need the first to get the second. The whole premise of catalog intelligence is to deliver the second on top of whatever you run today.




**07** Where this sits on the curve

## Agentic commerce is two stages ahead of where the market lives

Most agentic commerce content jumps straight to autonomous agents buying on a shopper's behalf. The market is not there. Commerce has moved through a long sequence of stages, from paper catalogs to digital PDFs to early portals to AI-enabled chat, and only recently toward agents.



The gap is the point. A company on a PDF cannot leap straight to autonomous agents. But it does not have to climb every stage between either. The layer is what lets a business skip the rebuild and still be ready for the surfaces that are coming.



08

## What to do now

### A path that is not a migration

01

#### Find where the product knowledge lives

Inventory the real sources. PDF catalogs, spreadsheets, rep knowledge, the fields buried in your ERP. Most of what an agent needs already exists. It is just scattered.

02

#### Get it structured

Turn those sources into machine-readable attributes, pricing, and order rules. This is the actual data-foundation work, and it can be done from what you have without replacing anything.

03

#### Make it answerable

Expose the structured catalog through a buying experience people can use by text, voice, or image, and through endpoints an agent can call. One is for today's buyers. The other is for the agents arriving next.

04

#### Measure the leak you close

After-hours demand, abandoned searches, quotes that never got a reply. These are the dollars the current experience loses every day. They are also the cleanest way to prove the layer paid for itself.

**None of these steps replaces your website or your ERP.** Each one moves the catalog closer to agent-ready while the existing business keeps running and selling.



09 The bottom line

## The companies that make their catalog answerable will be selling while the rest scope a migration

The data foundation argument is right. Agentic commerce will be won and lost on catalog data. The mistake is assuming the only way to fix the data is to rebuild the system that holds it.

For B2B, the faster and cheaper move is a layer that makes the catalog you already have legible to the agents that are coming. Do that, and you are answering buyers, and agents, in days. Skip it, and you are scoping a project the market will not wait for.

SEE IT ON YOUR OWN CATALOG

### Your catalog should be selling. Even when you're not.

ChatSKU ingests your existing catalog, including a PDF, and turns it into a live buying experience buyers can search, quote, and order by text, voice, or image. One line of embed code. No migration, no ERP rebuild, no IT project.

[Start My Free Trial](#)

[See How It Works](#)

Pay on your first quote.



Gigi JK

Gigi JK is the founder of ChatSKU and Virtina, bringing more than 28 years of experience across digital transformation, eCommerce strategy, AI-driven growth systems, and business modernization. His work spans startups, scale-ups, and SMBs, with a focus on turning complex operational problems into practical growth frameworks. Before ChatSKU, Gigi built and scaled a seven-figure eCommerce business and led Virtina as an eCommerce engineering and business transformation consultancy. At ChatSKU, he focuses on helping B2B manufacturers, distributors, and wholesalers make complex catalogs searchable, quote-ready, and agent-ready without forcing a full platform rebuild.

Sources. Gartner B2B buyer surveys (2026): 67% prefer a rep-free path, 70% want fully self-service. Sana Commerce B2B Buyer Report (2025): 75% would switch suppliers for a better experience, 91% among US buyers, 41% of manufacturing buyers cannot locate products on existing stores. LinkedIn, The Trust Advantage (2026). Maturity-stage estimates from the ChatSKU B2B Commerce Evolution framework.

Note. Figures are directional, drawn from published industry research, and provided to frame the decision rather than to guarantee a result.