



CommerceStudio.ai
ARCHITECTING B2B COMMERCE

A COMMERCESTUDIO.AI FIELD GUIDE

The Operability Shift

A framework and self-diagnostic for B2B commerce operators reading the new constraint correctly.

VERSION

1.0

PUBLISHED

2026

FORMAT

Framework & Diagnostic

The Operability Shift Field Guide

Version 1.0 · Published 2026

CommerceStudio.ai is a B2B commerce advisory practice working with operators, platforms, and integrators across the B2B commerce ecosystem. The Operability Shift framework and Diagnostic described in this Field Guide are CommerceStudio.ai intellectual property, developed from observations across our advisory work.

This Field Guide may be shared in its complete, unmodified form for educational and internal-strategy purposes. The framework, diagnostic methodology, and maturity tier definitions may be referenced with attribution to CommerceStudio.ai.

For advisory engagements, the Operability Diagnostic structured assessment, or to subscribe to The CommerceStudio.ai Dispatch, visit commercestudio.ai.

© 2026 CommerceStudio.ai · Architecting B2B Commerce

The Operability Shift Field Guide

FRONT MATTER

Executive Summary	4
-------------------------	---

PART ONE · THE FRAMEWORK

The Constraint Has Moved	7
Signal 1: Discovery Has Moved Off the Storefront	9
Signal 2: The Operating Model Has Become the Constraint	11
Signal 3: AI Value Lands First in Existing Work	13
The Framework at a Glance	15

PART TWO · THE DIAGNOSTIC

How the Diagnostic Works	17
The Twelve Dimensions	18
Scoring Worksheet	21
The Four Maturity Tiers	22

PART THREE · APPLICATION

What to Do With Your Tier Result	24
Communicating Findings Internally	25
About CommerceStudio.ai	26

The B2B commerce constraint has moved, and most teams haven't caught up.

For the last five years, B2B commerce conversations have centered on architecture. The teams we advise have already made those decisions. What's slowing them down now is the ability to operate what they've built.

The Operability Shift is the pattern we have observed across B2B operators over the past 18 months: the binding constraint in B2B commerce has moved from platform architecture to operating model. Discovery is migrating to surfaces operators don't control. Operating velocity is bottlenecked not by tools but by alignment, workflow, and decision cycles. AI investment is producing returns when it compresses existing work and underperforming when it chases new buyer experiences.

This Field Guide codifies the shift into a usable framework, a structured self-diagnostic, and a set of recommendations for B2B operators reading the change correctly. The framework identifies three signals. The diagnostic evaluates twelve dimensions across those signals and produces a maturity tier. The application section outlines what to do based on where you sit.

Who this Field Guide is for

Commerce leaders, IT leaders, and operating executives at B2B manufacturers, distributors, and platform-served businesses. The framework applies whether you serve B2B only, B2C only, or both. The diagnostic is designed to run in under thirty minutes and produce a tier and a prioritized next move.

The platforms are flexible. The teams aren't fast. The bottleneck has moved.

THE OPERABILITY SHIFT · SIGNAL TWO

The three signals of the Operability Shift, in brief:

- **Signal 1 — Discovery has moved off the storefront.** Buyers are finding products through AI agents, LLMs, marketplaces, and AI assistants. The storefront is one channel among many, and product data must be agent-ready to participate in the buying journey.
- **Signal 2 — The operating model has become the constraint.** Platform flexibility no longer translates to operating velocity. The bottleneck is cross-functional alignment, workflow integration, and decision cycle time.
- **Signal 3 — AI value lands first in existing work.** The fastest B2B AI ROI comes from compressing known operational workflows. AI investment in new buyer experiences produces returns later and is less consequential than the operational compression layer.

PART ONE

The Framework

The shift from architecture to operability, the three signals that define it, and the evidence behind each.

01

The Constraint Has Moved

For most of the last decade, B2B commerce conversations have centered on platform architecture. Headless or monolith. Composable or suite. The assumption was that picking the right stack would unlock the next phase of growth. That assumption is no longer holding.

The B2B operators we advise have already made architectural decisions. Their platforms are modern. Their stacks are composable. Their integrations are in place. And yet, in most cases, the speed they expected hasn't materialized. New experiences take quarters to ship. Cross-team alignment slows everything down. AI initiatives stall in pilot.

The architectural conversation that drove the last five years of B2B commerce investment was real and necessary. Most B2B businesses needed to modernize. But that work is now substantially complete for a meaningful portion of the market, and the conversation has not yet caught up to where the binding constraint has moved.

What changed

Three forces converged to shift the B2B commerce constraint from architecture to operability.

Architectural modernization reached saturation. A wave of B2B replatforming over 2019 to 2024 moved most ambitious B2B operators onto modern platforms. The marginal value of more architectural flexibility decreased as flexibility became table stakes.

Discovery migrated off the storefront. AI assistants, LLM-driven search, retail media, marketplaces, and feed-distributed surfaces became significant share-of-discovery channels. The catalog now has to perform on surfaces the operator does not control.

AI made operational work more compressible. Manual workflows that consumed significant operating capacity, particularly in B2B order processing, approvals, reconciliation, and catalog management, became candidates for AI-driven compression. The shape of B2B AI ROI became visible.

Three forces, one shift

The Operability Shift is what happens when these three forces land together. The work to be done has moved from platform selection and architecture to data work, operating model design, and AI investment placement. The framework that follows organizes those changes into three signals.

A note on definitions

Operability, as we use the term, refers to the organizational capacity to operate a commerce platform at the speed the business requires. It includes data quality and distribution, workflow integration, cross-functional decision velocity, and the placement of AI investment against operating priorities. It is distinct from **architecture**, which refers to the platform and stack choices that determine what is technically possible.

Signal One: Discovery Has Moved Off the Storefront

B2B buyers are increasingly finding products through AI agents, LLM-driven search, syndicated marketplaces, and AI assistants embedded in productivity software. The storefront is one channel among many in a journey that often converges elsewhere.

For most of the storefront's history, the storefront has been treated as the destination. Marketing, merchandising, content, search experience, and conversion optimization all assumed the storefront was where the buying journey converged. That assumption is breaking down.

Conference panels, analyst reports, and platform announcements over the past 12 months have all converged on the same pattern: a meaningful and growing share of B2B discovery happens before the buyer ever reaches the storefront. Buyers arrive pre-qualified by AI systems that have already evaluated dozens of options.

Whether your catalog appeared in those evaluations, and whether your business participated in the buying journey at all, depends on how agent-ready your product data is.

What an agent sees

When an AI agent operates against product data, it is not browsing a storefront. It is reading structured data from many sources, synthesizing across them, applying the buyer's stated and inferred intent, and producing a shortlist in seconds.

The data the agent reads is mostly not the marketing copy on the product detail page. It is the structured product feed flowing to AI surfaces through systems like Google's Universal Commerce Protocol, retail media platforms, payment-network commerce surfaces, or marketplace integrations. The agent reads attributes, taxonomy, identifiers, relationships, and availability signals. If those signals are missing, incomplete, or ambiguous, the product does not surface, and the operator usually never knows it didn't.

The four characteristics of agent-ready B2B product data

Structured

Specifications, dimensions, technical attributes, certifications, materials, and tolerances live in machine-readable fields, not in PDF datasheets or product image overlays.

Cross-Referenced

Compatible parts, alternative SKUs, complementary products, applications, and customer part numbers are modeled as explicit relationships.

Identified

Rich identifier systems (manufacturer part numbers, GTINs, customer part numbers) are present, consistent, and cross-linked across catalog, ERP, and customer-facing systems.

Live

Real-time signals on inventory, pricing intent, and lead times are exposed in the catalog and distributable to external surfaces.

Why this matters more in B2B

A B2C agent recommending a product mostly needs to surface the product itself. A B2B agent often needs to answer "what works with this," "what's the equivalent if this is out of stock," "what's the upgrade path," or "what application is this used for." Those answers require relational data that most B2B catalogs don't maintain explicitly.

This is also where customer part numbers become strategic. Most B2B buyers think in their own internal part numbers, not the supplier's. When an agent processes a buyer query that references a customer part number, the catalog has to cross-reference that to the actual SKU, or the product is functionally invisible.

Your catalog used to be optimized to convert humans who arrived on your storefront. It now has to be optimized to be readable by machines that decide whether humans arrive at all.

Signal Two: The Operating Model Has Become the Constraint

Platforms are flexible. Teams aren't fast. The binding constraint in B2B commerce has moved from technology choice to operating model maturity.

Across the operators we advise, technology choices have largely been made. Stacks are modern. Platforms are capable. Integration patterns are in place. The friction that is slowing teams down is no longer architectural. It is operational.

The patterns we see consistently

- Complex internal stakeholder structures that delay decisions.
- Slow buying and approval cycles that stretch project timelines.
- Disconnected workflows between marketing, merchandising, sales, and IT.
- The persistent challenge of cross-functional alignment around shipping change.

In that environment, additional architectural flexibility does not translate into additional speed. It often adds more to manage. Operators that have invested in flexibility without investing in operating model design end up sitting on platforms they cannot fully use.

When flexibility outruns the operating model

The multi-segment dimension

For B2B operators serving multiple segments, the operating model constraint compounds. Manufacturers running direct-to-consumer pilots, distributors launching consumer brands, B2B-first companies extending into D2C, and D2C brands adding wholesale all face an additional layer of friction.

Most of these merchants have had three bad options: pick one platform and hack the other use case onto it, run two platforms with duplicated effort, or build a Frankenstein. None scale. All produce operating tax that compounds over time.

Operators that have done the operating model work for multi-segment business operate at a meaningfully different pace than those who haven't. This is one of the highest-leverage operating model investments available to B2B businesses today.

The uncomfortable implication

Many teams that replatformed in the last two years are now sitting on more flexibility than they can use. The next phase of work for those organizations is not another platform decision. It is the operating model work that turns flexibility into pace: workflow integration, decision velocity, cross-functional alignment, and multi-segment coherence.

A quick test for Signal 2 maturity

If a new pricing change in your business requires coordination between four or more teams, takes longer than two weeks to ship, or routinely produces inconsistent results across customer-facing surfaces, your operating model is the constraint. Architecture and tooling are not the bottleneck.

Signal Three: AI Value Lands First in Existing Work

The fastest B2B AI ROI comes from compressing known workflows. AI investment in new buyer experiences produces returns later and is less consequential than the operational compression layer.

Across our advisory work, the AI use cases that get funded and produce ROI in B2B share a common shape. They compress known workflows. They remove manual steps. They make repeatable work faster. They do not invent new buyer experiences before fixing the operational workflows underneath them.

Where operational AI value lands first

The B2B workflows that benefit from AI compression first share three traits: they are repetitive, they are load-bearing for downstream activity, and they currently consume significant manual capacity.

- **Purchase order processing.** Manual PO intake, parsing, validation, and order creation. AI removes the manual steps without changing the workflow conceptually.
- **Invoice and payment reconciliation.** Matching invoices to POs to receipts to payments. High-volume, rule-based, and historically slow.
- **Approval routing.** B2B approval workflows with budget rules, role rules, and product rules. AI handles the routine routing and surfaces only the exceptions.
- **Catalog enrichment.** Extracting structured data from PDFs, classifying products, applying taxonomy. The same work that supports Signal 1 agent-readiness.
- **Customer service deflection.** Answering questions buyers ask repeatedly about availability, lead times, account status, and order history.
- **Returns and RMA triage.** Classifying return requests, applying policy, routing to disposition.

Why the new buyer experience narrative is misleading

Most B2B AI conversations are focused on new buyer experiences: conversational commerce, agentic shopping assistants, AI-driven product recommendations on the storefront. These will eventually matter. But they are not where B2B AI investment pays back first, and the operators chasing them before fixing the operational workflows underneath are spending budget in the wrong order.

B2B buyers, broadly, are not asking for new ways to interact with suppliers. They are trying to complete known tasks more efficiently. The fastest ROI from AI in B2B comes from making those tasks faster, not from inventing new ones.



The fastest ROI from AI in B2B comes from compressing the workflows everyone already understands. New buyer experiences land second.

THE OPERABILITY SHIFT · SIGNAL THREE

The Operability Shift Framework

Three signals. Twelve dimensions. One pattern reshaping B2B commerce.

01 Discovery Has Moved Off the Storefront

Buyers find products through AI agents, LLMs, marketplaces, and AI assistants. Visibility in those journeys depends on agent-ready product data: **Structured, Cross-Referenced, Identified, Live.**

02 The Operating Model Has Become the Constraint

Platforms are flexible. Teams aren't fast. The bottleneck is workflow, alignment, decision velocity, and multi-segment coherence. Architecture has been done; operating model is the next investment frontier.

03 AI Value Lands First in Existing Work

The fastest B2B AI ROI comes from compressing known workflows: POs, approvals, reconciliation, catalog enrichment, service deflection. New buyer experiences land second.

Reading the signals together

The three signals are not independent. They compound. Strong agent-readiness with a weak operating model produces visible discovery and invisible execution. Strong operating model with AI investment misplaced into new experiences produces velocity in the wrong direction. The diagnostic that follows evaluates all three signals together and surfaces the lagging one as the priority next move.

PART TWO

The Diagnostic

A structured self-assessment that evaluates twelve dimensions across the three signals and produces a maturity tier.

02

PART 2 · METHODOLOGY

How the Diagnostic Works

Twelve questions. Four dimensions per signal. Each scored on a 0–3 scale. The total score maps to one of four maturity tiers.

The Operability Diagnostic evaluates each of the three signals across four dimensions, for twelve dimensions total. The dimensions are designed to be self-assessable in under thirty minutes by an operator with leadership-level visibility into commerce, data, integration, and AI investment in their business.

Scoring scale

Each dimension is scored on a 0 to 3 scale.

- 0 **Not yet started** or addressed
- 1 **Partially in place**, with significant gaps
- 2 **Mostly addressed**, with known limitations
- 3 **Fully addressed**, comprehensive coverage

Total possible score is 36. The score maps to one of four maturity tiers, each with a different implied next move.

How to score honestly

The diagnostic is only as useful as the honesty of the scoring. We recommend three practices:

- 1 **Score against current state, not aspirational state.** "We plan to address this" is a 0 or 1, not a 2.
- 2 **Score against the full scope of the business.** If your catalog is agent-ready for the top 200 SKUs but not the long tail, the dimension is partial (1), not mostly addressed (2).
- 3 **Involve more than one stakeholder.** Commerce, IT, and operations will see different things. The dimensions that one role rates highly and another rates poorly are usually the most consequential ones to discuss.

PART 2 · THE TWELVE DIMENSIONS

Signal 1 · Discovery Readiness

Four dimensions evaluating whether your product catalog is positioned for AI-mediated discovery.

1.1 Structured

How much of your product specification data lives in structured, queryable fields rather than PDFs, images, or unstructured descriptions?

- 00 ☐ Most spec data sits in PDFs or images
- 01 ☐ Some core specs are structured; significant gaps remain
- 02 ☐ Most specs are structured, with some legacy PDFs
- 03 ☐ All product specs are fully structured and queryable

1.2 Cross-Referenced

How systematically are product relationships (alternatives, compatibility, applications, customer part numbers) modeled in your catalog?

- 00 ☐ Relationships are not modeled; salespeople hold this knowledge
- 01 ☐ Limited relationships exist for core SKUs
- 02 ☐ Most products have explicit relationships; gaps in the long tail
- 03 ☐ Comprehensive relationship modeling across the catalog

1.3 Identified

How consistent are your product identifiers (manufacturer part numbers, GTINs, customer part numbers) across catalog, ERP, and customer-facing systems?

- 00 ☐ Significant identifier inconsistency exists
- 01 ☐ Some identifier management; known gaps and conflicts
- 02 ☐ Identifiers are mostly consistent across systems
- 03 ☐ Identifier hygiene is fully managed and cross-linked

1.4 Live

Are real-time signals (inventory, pricing intent, lead times) exposed in your catalog and distributable to external surfaces?

- 00 ☐ Real-time signals are not exposed externally
- 01 ☐ Some signals are available, inconsistently
- 02 ☐ Most signals are exposed, with known limitations
- 03 ☐ All key signals are real-time and externally distributable

PART 2 · THE TWELVE DIMENSIONS

Signal 2 · Operating Model Maturity

Four dimensions evaluating whether your organization can move at the speed your platform allows.

2.1 Cross-Functional Alignment

How quickly can your organization align commerce stakeholders (marketing, merchandising, IT, sales) to make a meaningful change?

- 00 ☐ Alignment routinely takes a quarter or more
- 01 ☐ Alignment takes weeks to months depending on scope
- 02 ☐ Most decisions ship within a few weeks
- 03 ☐ Alignment is fast and structured

2.2 Workflow Integration

How many manual handoffs exist between commerce systems (catalog, OMS, ERP, CRM, marketing automation)?

- 00 ☐ Many manual handoffs; significant operational tax
- 01 ☐ Several manual handoffs remain
- 02 ☐ Most workflows are integrated, with some manual touchpoints
- 03 ☐ Workflows are integrated end-to-end

2.3 Change Velocity

How quickly can a new commerce experience (page, feature, integration) go from concept to live?

- 00 ☐ Months
- 01 ☐ Weeks
- 02 ☐ Days
- 03 ☐ Hours for configuration; days for engineering work

2.4 Multi-Segment Coherence

How well does your operating model handle multi-segment complexity, or how prepared is your organization to add a new segment without doubling overhead?

- 00 ☐ Segments would require entirely separate implementations
- 01 ☐ A new segment would require significant duplicated processes
- 02 ☐ Some processes are unified; segment-specific handling layered on
- 03 ☐ Unified operating model that supports multiple segments cleanly

PART 2 · THE TWELVE DIMENSIONS

Signal 3 · AI Value Placement

Four dimensions evaluating whether your AI investment is concentrated where it produces operational ROI.

3.1 Workflow Compression Focus

How is your AI investment prioritized, toward compressing existing workflows or inventing new buyer experiences?

- 00 ☐ AI investment is concentrated in new buyer experience pilots
- 01 ☐ Some operational workflow AI; most focus on new experiences
- 02 ☐ Balanced operational and experience AI
- 03 ☐ AI investment prioritizes operational workflow compression

3.2 Operational AI in Production

How much of your AI investment is in production today versus in pilots?

- 00 ☐ Mostly pilots, limited production deployment
- 01 ☐ Some production AI; most still in pilot stage
- 02 ☐ Multiple AI capabilities in production
- 03 ☐ AI is integrated across operational workflows in production

3.3 Measurement Framework

Do you have a clear measurement framework for AI value (time saved, error reduction, throughput)?

- 00 ☐ AI value is not systematically measured
- 01 ☐ Some measurement exists, mostly anecdotal
- 02 ☐ Structured measurement on core AI investments
- 03 ☐ Comprehensive measurement framework across AI initiatives

3.4 AI Roadmap Alignment

Is your AI roadmap aligned to operating model priorities, or driven by available tools?

- 00 ☐ AI roadmap is tool-driven
- 01 ☐ Roadmap is partly aligned, partly tool-driven
- 02 ☐ Roadmap is mostly aligned to operating priorities
- 03 ☐ AI roadmap is explicitly designed to compound operating velocity

PART 2 · SCORING

Scoring Worksheet

Transfer your scores from each dimension. Sum by signal. Read your tier on the next page.

DIMENSION	SIGNAL TOTAL	YOUR SCORE (0-3)
1.1 Structured		
1.2 Cross-Referenced		
1.3 Identified		
1.4 Live	Signal 1 sum	
2.1 Cross-Functional Alignment		
2.2 Workflow Integration		
2.3 Change Velocity		
2.4 Multi-Segment Coherence	Signal 2 sum	
3.1 Workflow Compression Focus		
3.2 Operational AI in Production		
3.3 Measurement Framework		
3.4 AI Roadmap Alignment	Signal 3 sum	
Total Operability Score		

Total possible score is 36. Carry your total to the next page to read your maturity tier.

PART 2 · MATURITY TIERS

The Four Maturity Tiers

Find your total score in the range column. Each tier carries a different implied next move.

TIER 1

Architecture Era

Score 0–9

You are operating in the architecture-era conversation. Platform choice is still treated as the binding constraint. Discovery is assumed to happen on the storefront. AI investment is in pilots or has not started.

Most Tier 1 organizations are not behind. They are working from the same playbook the industry used for the last decade. The shift to operability has happened, but the implications have not surfaced in their planning yet.

TIER 2

Mid-Shift

Score 10–18

You have recognized the shift and started moving. Some catalog investments are in motion. AI pilots have produced learning. The operating model has been discussed but not redesigned.

Organizations in this tier typically have unbalanced scores. One signal is advanced while the others lag. The risk is that the advanced signal looks like progress while the lagging signals compound into structural friction.

TIER 3

Operability Ready

Score 19–27

You have invested in operability across all three signals. Catalog data is largely agent-ready. Operating model decisions ship fast. AI investment is mostly directed at workflow compression.

Organizations at Tier 3 are positioned to compound. The next gains come from execution and refinement, not foundational change. Specific friction points become visible and addressable.

TIER 4

Compounding

Score 28–36

You are operating at velocity. The Operability Shift is not a future state for your organization. It is the operating reality. Discovery readiness is comprehensive. The operating model is fast. AI investment is producing measurable returns.

Organizations at Tier 4 are in a small minority. The capability gap between Tier 4 and Tier 1 is significant and growing.

PART THREE

Application

What to do with your tier result, how to communicate findings, and how to engage CommerceStudio.ai.

03

What to Do With Your Tier Result

The diagnostic produces a score and a tier. What matters is what you do with them.

Tier 1 (Architecture Era) — Bring operability into the executive conversation

Most Tier 1 organizations have not had the framing that lets them see the shift. The first move is internal. Surface the catalog work, the operating model work, and the AI prioritization work as first-class executive priorities, not as items downstream of platform selection. The diagnostic itself is a useful conversation starter for leadership.

Tier 2 (Mid-Shift) — Commit dedicated investment to your lowest-scoring signal

Tier 2 scores are typically uneven. The strongest signal can mask the lagging ones, and the lagging signals compound into structural friction. Identify your weakest signal and structure the next 6 to 12 months around closing it before the next signal becomes the limiter.

Tier 3 (Operability Ready) — Move to measurement and refinement

The structural work is done. The organizations that compound at Tier 3 are the ones that explicitly track operating velocity, discovery share, and AI workflow ROI. Without measurement, gains plateau. The work from here is sequencing the highest-ROI refinements and building the metric system that makes compounding visible.

Tier 4 (Compounding) — Lead the market conversation

You are in a small minority of organizations operating at velocity. The first move from Tier 4 is contribution. Documented case studies, framework refinements, and contributions to the broader market conversation are the leadership work that returns to the organization in talent, partnership, and category authority.

PART 3 · APPLICATION

Communicating Findings Internally

How to translate a diagnostic result into executive alignment and prioritized action.

Running the diagnostic produces a number and a tier. Turning those into action requires translating the result into the language and frame of reference your leadership team uses for capital allocation, roadmap planning, and operating model decisions.

Three communication patterns we recommend

Lead with the binding constraint, not the score. "Our binding constraint is operating model maturity, not platform capability" is more actionable than "We scored 14." The score supports the argument; it should not be the argument.

Frame the lagging signal as the priority, not the lowest score. "Our agent-readiness gap is the most consequential investment we are not making" lands differently than "We scored 2 out of 12 on Signal 1." The substance is the same; the framing carries different weight.

Tie tier movement to business outcomes. "Moving from Tier 2 to Tier 3 is what unlocks the operating velocity required to ship the new B2B program on time" connects diagnostic work to a real business deliverable. Tier movement without business framing is interesting but not actionable.

When to involve external advisory

Most operators can run the diagnostic and produce a tier internally. The places where external advisory adds the most value are: aligning leadership teams that are reading the shift differently, building the roadmap that translates tier movement into sequenced investment, and benchmarking against patterns seen across other operators. CommerceStudio.ai engages on all three.

ABOUT

About CommerceStudio.ai

A B2B commerce advisory practice working with operators, platforms, and integrators across the B2B commerce ecosystem.

CommerceStudio.ai advises B2B commerce operators on strategy, operating model design, and the placement of investment against the patterns reshaping B2B commerce. Our work spans manufacturers, distributors, and platform-served B2B businesses across discovery readiness, operating model maturity, and AI value placement.

The Operability Shift framework and Diagnostic described in this Field Guide are central to how we engage with operators. Most of our 2026 advisory work starts with the Diagnostic and moves into structured engagements that address the highest-leverage signal for each client.

Three ways to engage

Structured Operability Diagnostic engagement. Executive interviews, system-level observation, and detailed roadmap output across all twelve dimensions. Output is a prioritized investment plan and an execution-ready next quarter.

Advisory engagements. Leadership-level work on platform direction, operating model design, and roadmap sequencing against the Operability Shift framework. We work directly with commerce leaders, IT, and operations.

The CommerceStudio.ai Dispatch. Bi-weekly analysis on the Operability Shift, platform-level signals, and the operators executing at each tier. Future Dispatch issues will publish anonymized aggregate Diagnostic data as the cohort grows.

Engage CommerceStudio.ai

To discuss a Diagnostic engagement, advisory project, or to subscribe to The CommerceStudio.ai Dispatch, visit commercestudio.ai or contact us directly.

COMMERCESTUDIO.AI · ARCHITECTING B2B COMMERCE