



# THE AI CONTEXT READINESS CHECKLIST



**20 questions to ask before you trust an AI system with a business question.**



# THE AI CONTEXT READINESS CHECKLIST

## 20 questions to ask before you trust an AI system with a business question.

A more capable model cannot fix a business question it doesn't understand.

Research testing frontier AI models on enterprise-style data found that giving them explicit business semantics — rather than a database schema alone — improved accuracy by 17 to 23 percentage points. And a technically correct answer can still be the wrong one if it wasn't built for the person asking it.

This checklist is a practical companion to *\*Why AI Needs Context, Not Just Data\** (The Data Advantage, Issue 9).


It's designed for data, analytics and operations leaders who want a fast, honest read on how ready their AI systems really are — not at the model layer, but at the layer underneath it: the business meaning, ownership and role-awareness that determine whether an answer can be trusted.

Work through the four sections. Check what's genuinely true today, not what's on the roadmap.



# 01 SEMANTIC FOUNDATIONS

Does your AI understand what your data means?



Every core business metric (revenue, active customer, service level, and so on) has one documented, agreed definition.



That definition lives somewhere an AI system can actually read it – not only in an analyst's head or a slide deck.



Table and field names are mapped to the business language people actually use when they ask a question.



You know which source system is authoritative when two systems disagree.



Someone can explain, in one sentence, why this month's number and last month's number aren't directly comparable – if that's true.



# 02 CONTEXT OWNERSHIP

Who is responsible for keeping this accurate?



|                          |                                                                                                                        |
|--------------------------|------------------------------------------------------------------------------------------------------------------------|
| <input type="checkbox"/> | One function — not every project team — owns how business definitions and relationships are maintained.                |
| <input type="checkbox"/> | There's a known process for updating a definition when the business changes it.                                        |
| <input type="checkbox"/> | Context built for one AI use case can be reused by the next one, rather than rebuilt from scratch.                     |
| <input type="checkbox"/> | You could name the person or team accountable if an AI-generated answer turned out to be wrong because of bad context. |
| <input type="checkbox"/> | Context quality is reviewed on a schedule — not only after something has already broken.                               |





## 03 PERSONA & ROLE AWARENESS

Does the same question get the right answer for the right person?



|                          |                                                                                                                                              |
|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------------|
| <input type="checkbox"/> | You've mapped out who typically asks which kinds of questions, and what a useful answer looks like for each of them.                         |
| <input type="checkbox"/> | Permissions determine what a person can see – and you've separately considered what they actually need to see.                               |
| <input type="checkbox"/> | A frontline employee and an executive asking the same question would get answers at appropriately different levels of detail.                |
| <input type="checkbox"/> | Sensitive or commercially confidential context (margin, compensation, legal exposure) is excluded from answers by default, not by exception. |
| <input type="checkbox"/> | Someone has tested this with a real question from a real role – not just reviewed it in theory.                                              |



# 04 VALIDATION & GOVERNANCE

Can you trust the answer — and prove it if asked?



|                          |                                                                                                                                         |
|--------------------------|-----------------------------------------------------------------------------------------------------------------------------------------|
| <input type="checkbox"/> | You can trace which sources and definitions contributed to a specific AI-generated answer.                                              |
| <input type="checkbox"/> | AI-assembled context is tested for accuracy the way you'd test a data pipeline — not just checked once at launch.                       |
| <input type="checkbox"/> | There's a way to flag and correct a wrong or outdated definition before it spreads to more answers.                                     |
| <input type="checkbox"/> | You know your obligations under relevant AI transparency rules (for example, EU AI Act Article 50) for how these answers are disclosed. |
| <input type="checkbox"/> | A non-technical stakeholder could look at your context architecture and understand, roughly, how an answer gets made.                   |



# WHERE DOES THAT LEAVE YOU?

Count your checked boxes out of 20.

**0–7**

## **Early stage**

Context is still living in people's heads and spreadsheets, not your systems. AI answers will be inconsistent until the semantic and ownership gaps close.

**8–14**

## **Developing**


The foundations exist, but ownership, persona-awareness or testing are inconsistent. This is where most organisations sit in 2026 — the gap is usually process, not technology.

**15–20**

## **Advanced**

You're positioned to scale AI reliably across the business, with context treated as shared infrastructure rather than a per-project afterthought.





## WANT HELP CLOSING THE GAPS?

See how AskEmite is built on a governed context layer  
· [emite.com/ask-emite](https://emite.com/ask-emite)

Talk to us · [emite.com/request-a-demo](https://emite.com/request-a-demo)



[www.emite.com](https://www.emite.com)

