

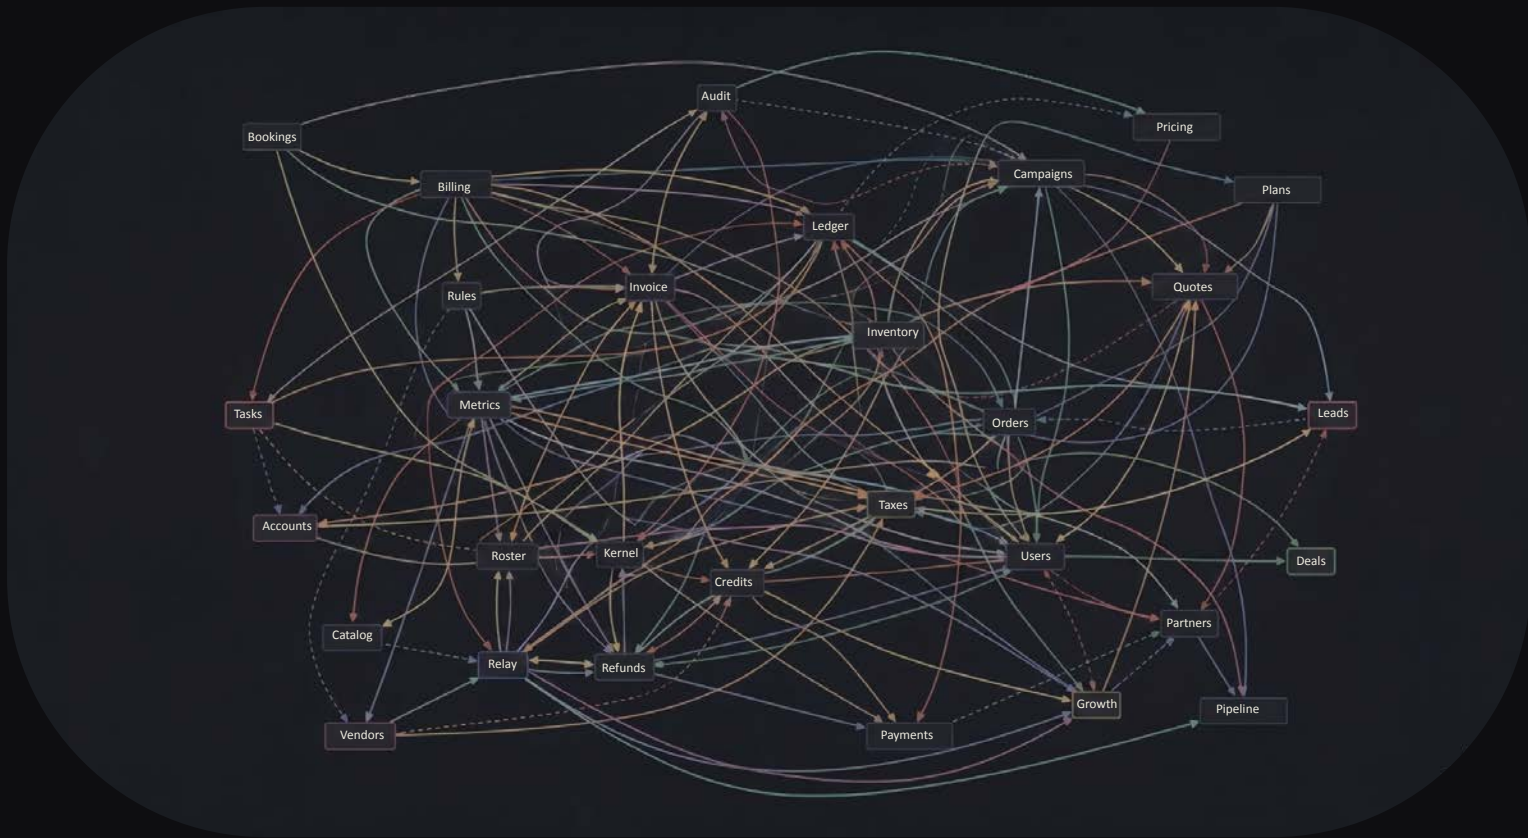
The Architecture Hub

Teaching AI to Understand Your System (Before It Codes It)

Lior Schejter | May 2026

How well does your AI agent understand
your **system**?

Not just one repo, but how 100+ services fit
together?

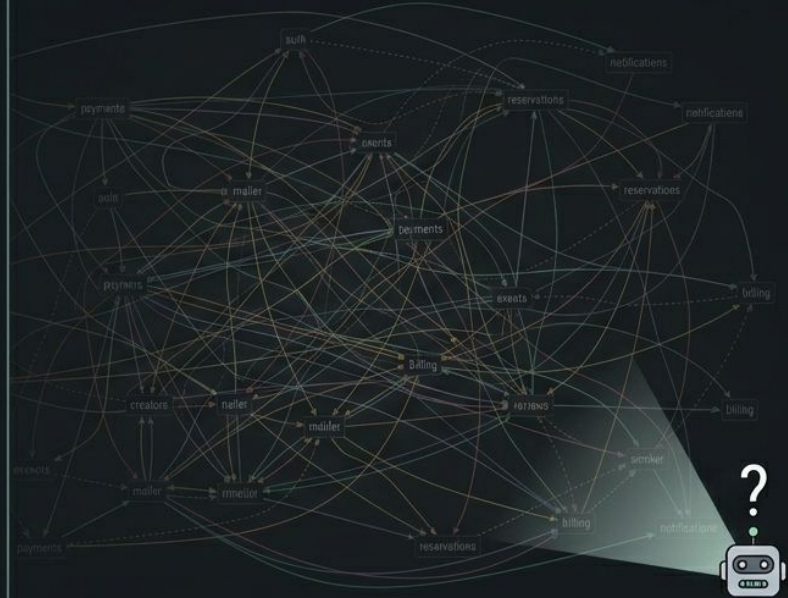


What AI Sees



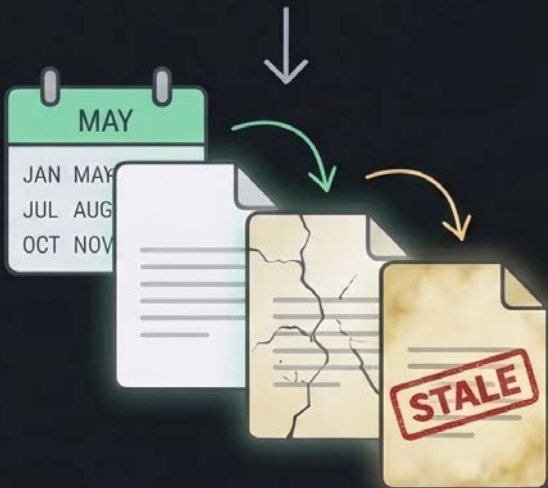
One Repository

What AI Needs to See





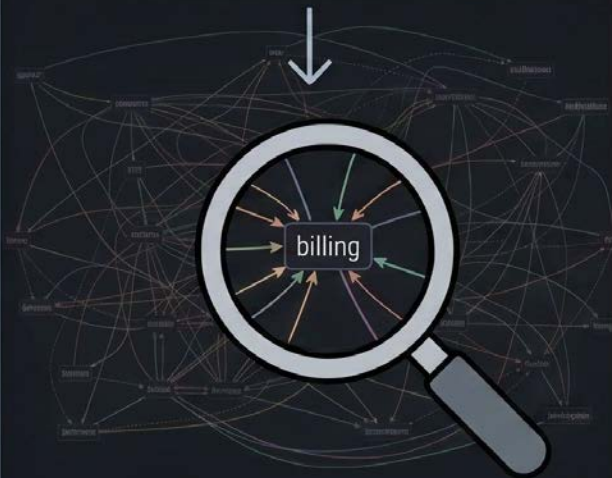
Human Documentation



Entropy wins.



Code Indexing



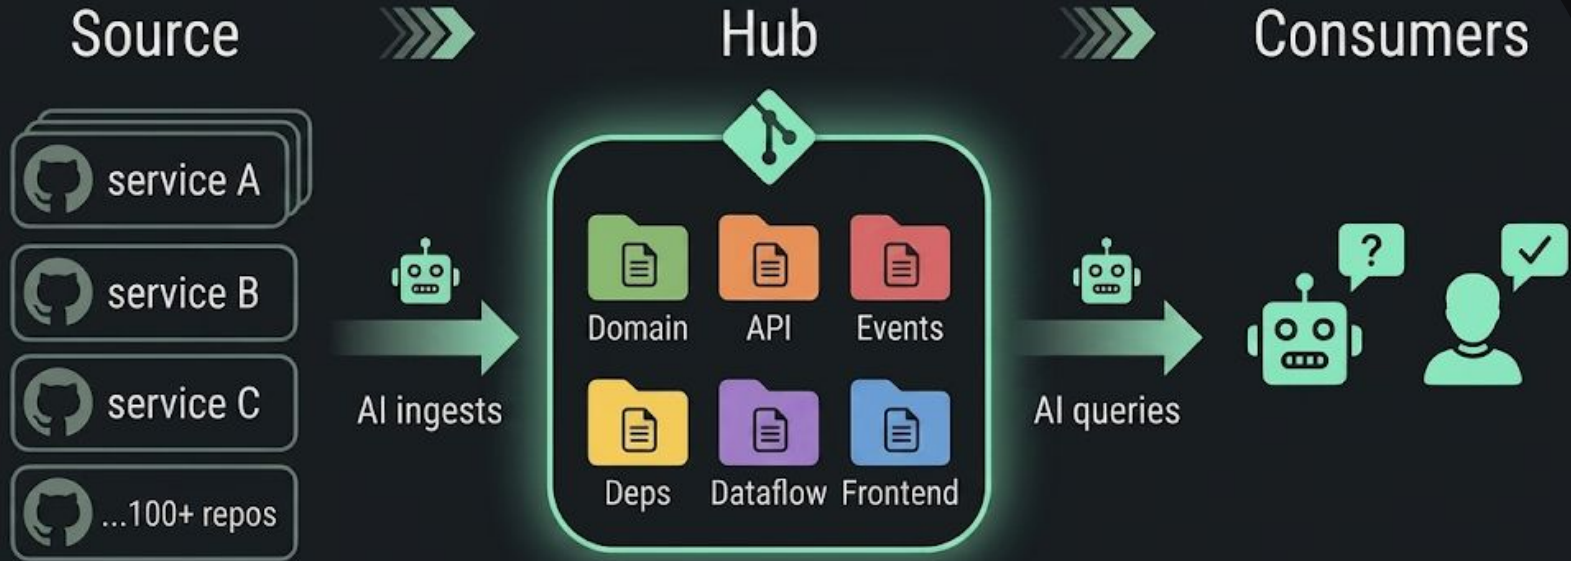
Single-repo vision.

AI is Smart Enough



... and Tireless

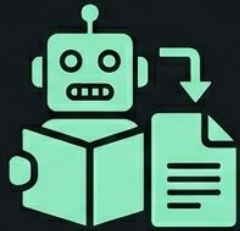
The Architecture Hub



Core Design Principles/Constraints



Zero
infrastructure



AI-native
I/O



Diffable &
reviewable



Humans in
the loop



Scales to
100+ repos

Files Beat Databases



Graph DB



Wiki



Embedded docs

VS.



Git + Markdown



Versioned
for free




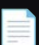



Zero-ops



LLM-friendly
format

Decompose by Perspective, Not Service

	Domain	API	Events	Deps	Dataflow	Frontend
reservations						
payments						
inventory						
orders						
clients						
...						

Understanding a service?

Following a flow?

Path are Computable, Not Discoverable

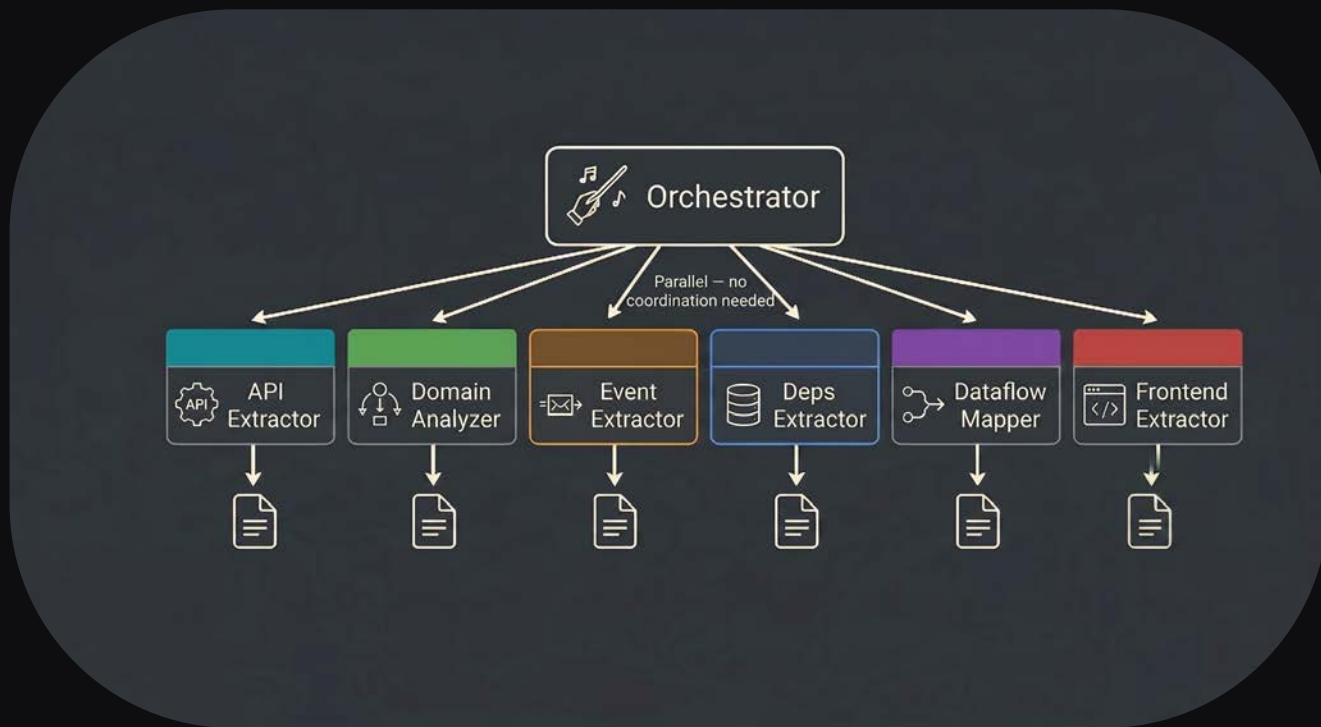
$f(\text{facet}, \text{repo}) \rightarrow \{\text{facet}\} / \{\text{repo}\} .\text{md}$

$(\text{api}, \text{payments}) \rightarrow \text{api} / \text{payments} .\text{md}$

$(*, \text{inventory}) \rightarrow * / \text{inventory} .\text{md}$

$(\text{events}, *) \rightarrow \text{events} / * .\text{md}$

Ingestion by Orchestrator + Subagents



Reasoning → AI

Verification → Deterministic Code

Before

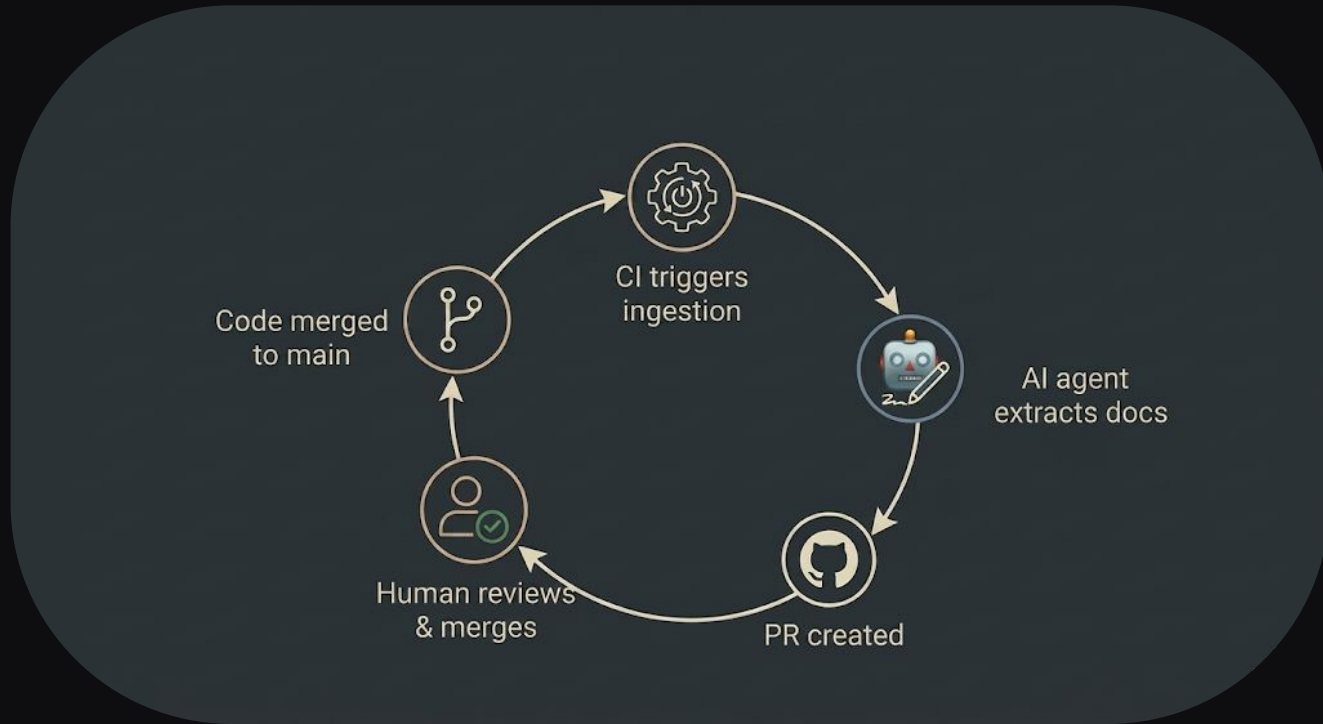


After

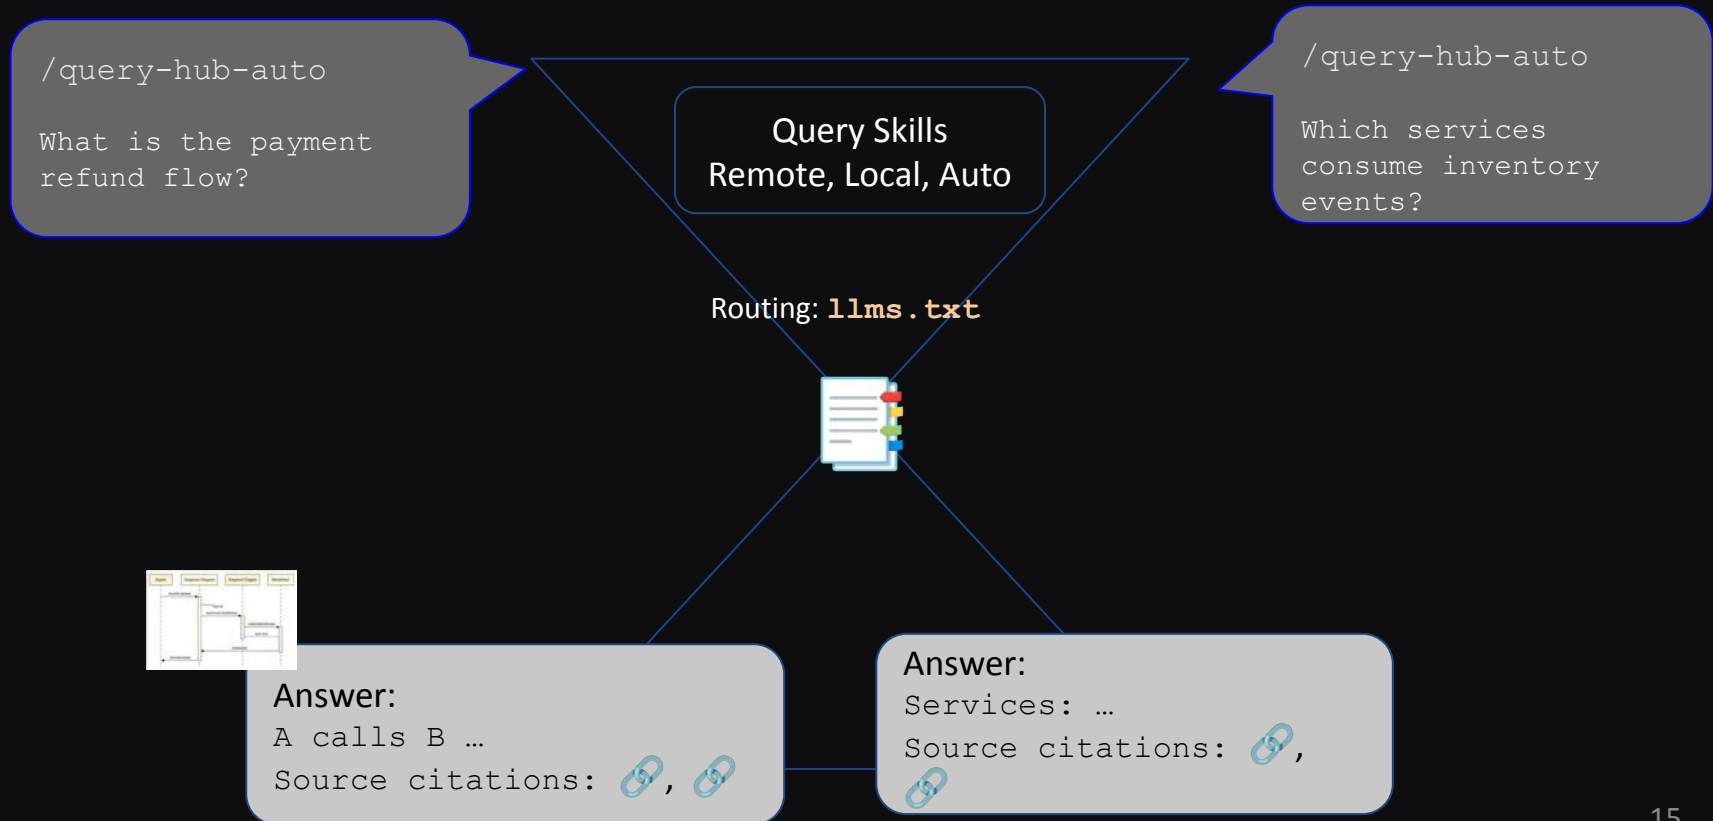


💡 Don't use AI where you don't need AI

Automate the Loop, Keep Humans Reviewing



Querying: Two Audiences, Same Substrate



Challenge: Ingestion Variance

Same repository
ingested twice



Different
documents



Structure
enforced



Grounded
in code

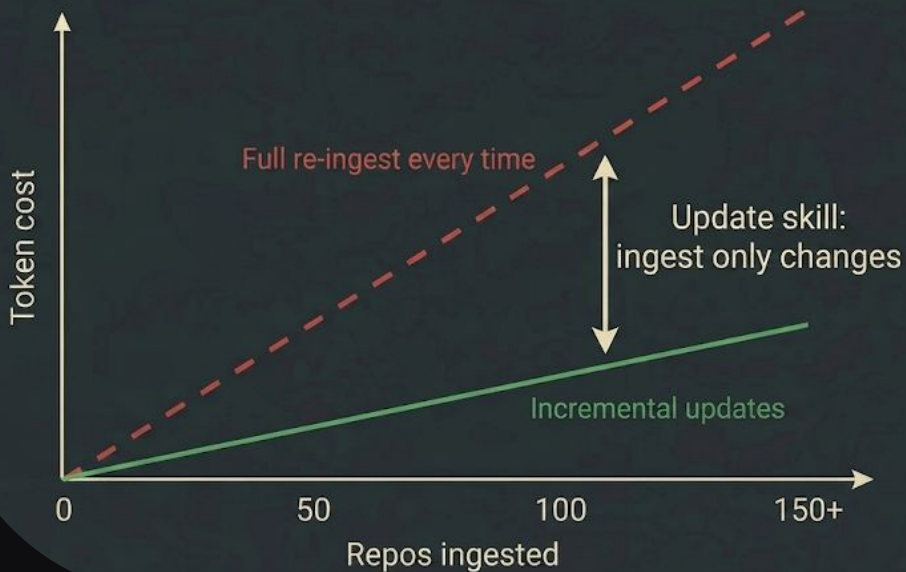


Mechanically
verified



Semantically
reviewed

Challenge: Cost and Scale



Merge to main

Triggers are selective

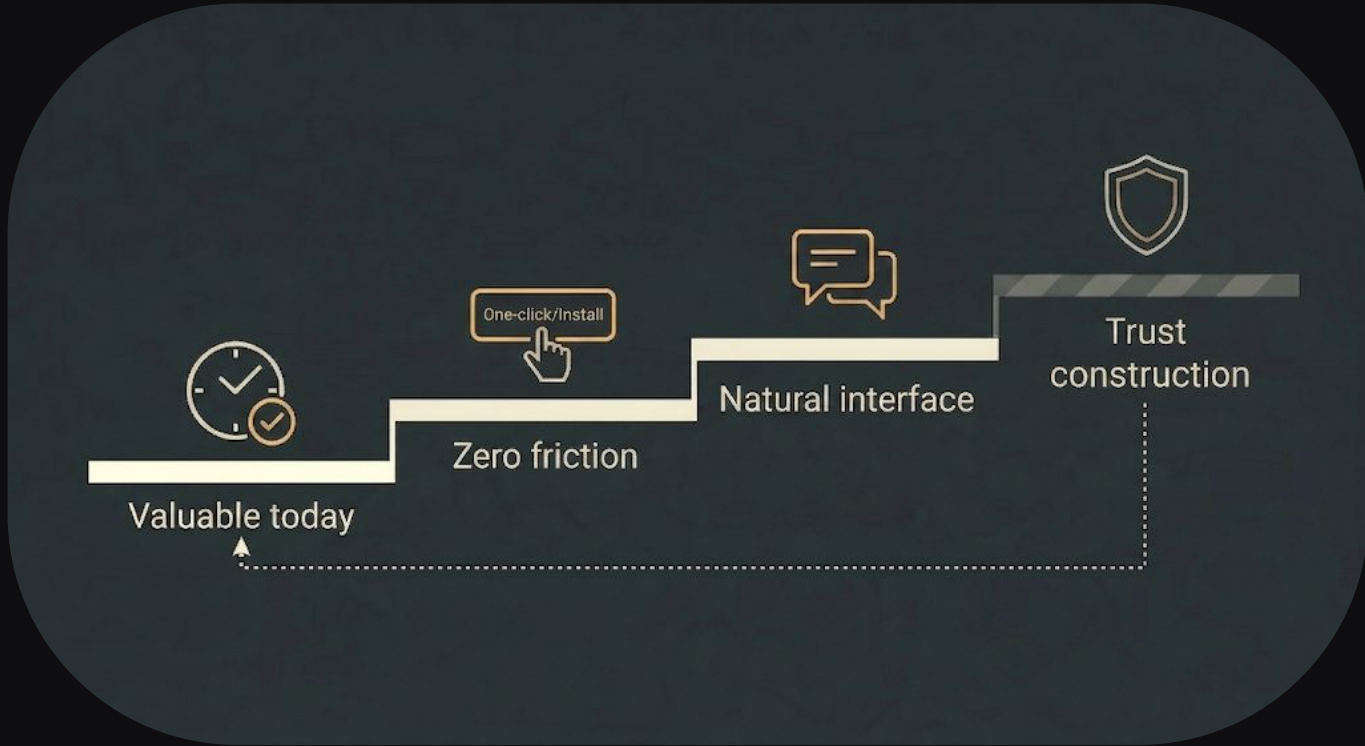


Re-ingest one facet,
not all six

1→2→3

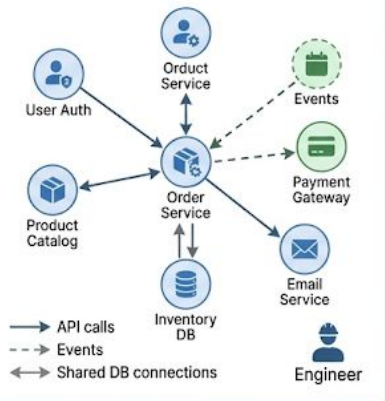
Sequential batch avoids
rate-limit cascades

Challenge: Adoption



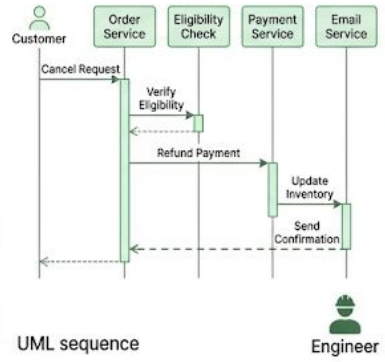
Structured Reports

1 Dependency Map



Blast radius in minutes

2 Flow Analysis



End-to-end trace with citations

3 Plain English Explainer

When a customer cancels an order, the system first checks eligibility, then reverses payment charges, releases held inventory, and sends a confirmation email...

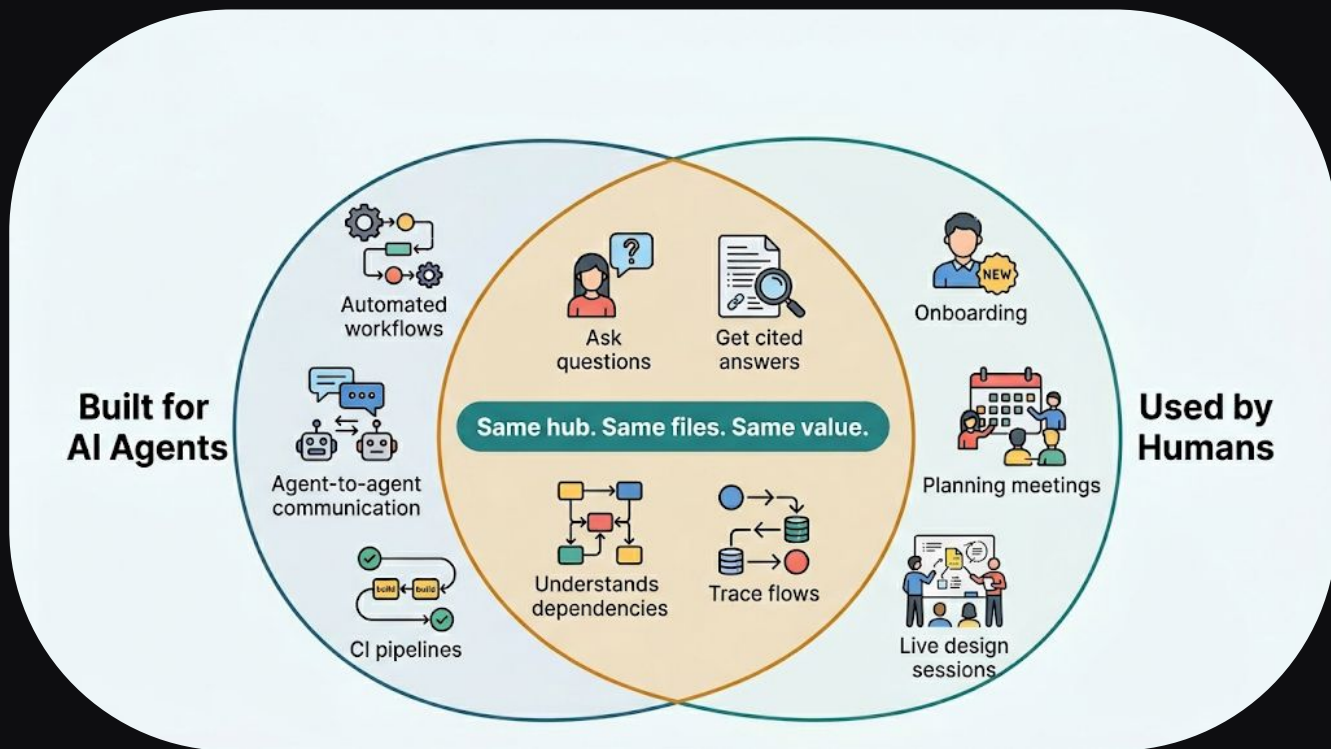
PM/businessperson

Same data, different audience

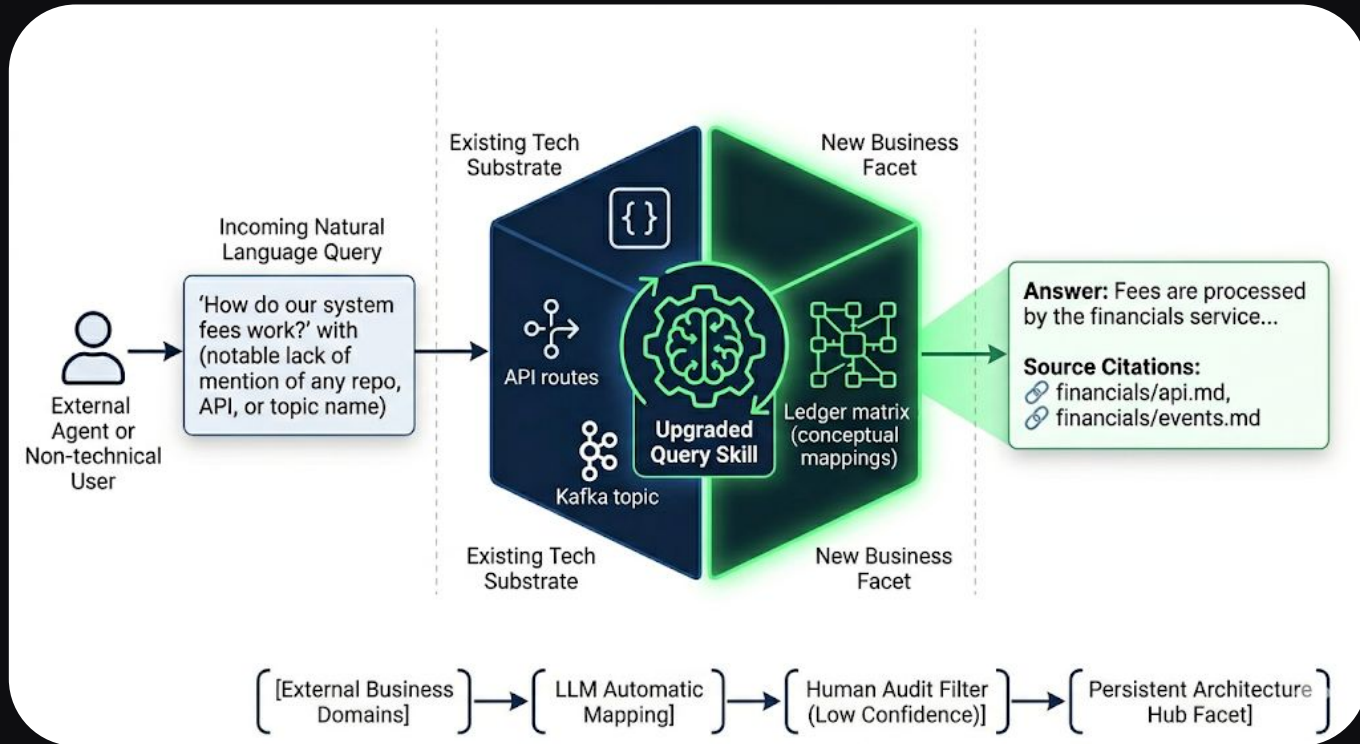
Unified Data Source

same substrate, three different outputs

Built For Humans and Agents





Mapping to Business Language





Key Takeaways


 Start with flat files. Git + Markdown.


 Decompose by facets/aspects

 Make indexing simple and paths computable

 AI for reasoning, scripts for verification

 Automate for production, gate with human review

 Citations are architecture, not decoration

 The value is in the query, not the document

 Humans are important consumers as well

Scale So Far

700+ 

markdown files

150 

repositories documented

6 

architecture facets

37 

agent skills

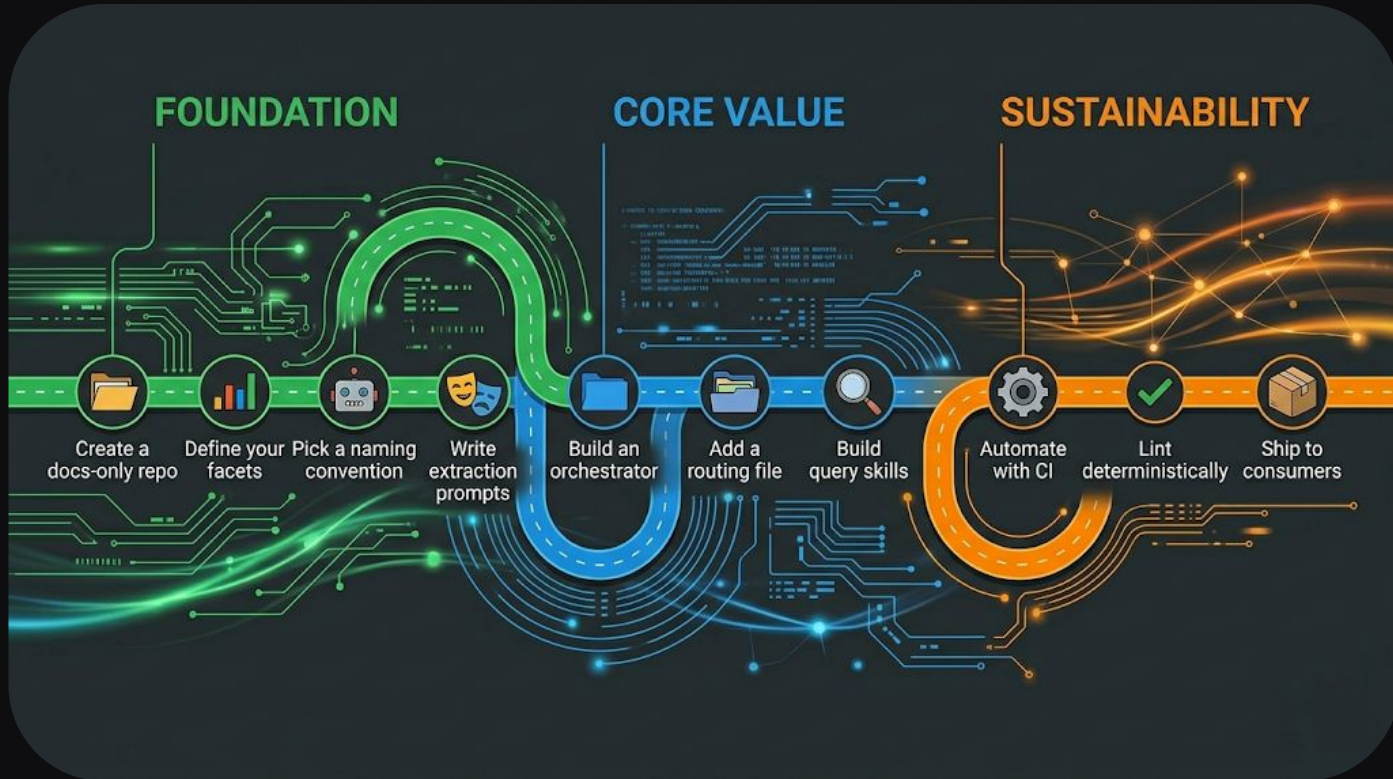
5 

report templates

~2.5 
months

from idea to production

Implementation Roadmap



Thank You!



[Bits and Bytes](#)



[The Architecture
Hub - Blog Post](#)



[Lior Schejter](#)