MASTER CLASS

Most Companies Don't Have an AI Problem. They Have a Context Problem

DIGIMARCON CANADA 2026
DIGITAL MARKETING CONFERENCE



John Bush

CO-FOUNDER AND PRESIDENT
REVENUE EXPERTS, INC

TORONTO, ON ~ APRIL 15 - 16, 2026
DIGIMARCONCANADA.CA | #DigiMarConCanada

2026 WORLD TOUR

WALLPAPERSWIDE.COM

DiGiMarCon Toronto | April 15th, 2026

MOST COMPANIES DON'T HAVE AN AI PROBLEM: THEY HAVE A CONTEXT PROBLEM.

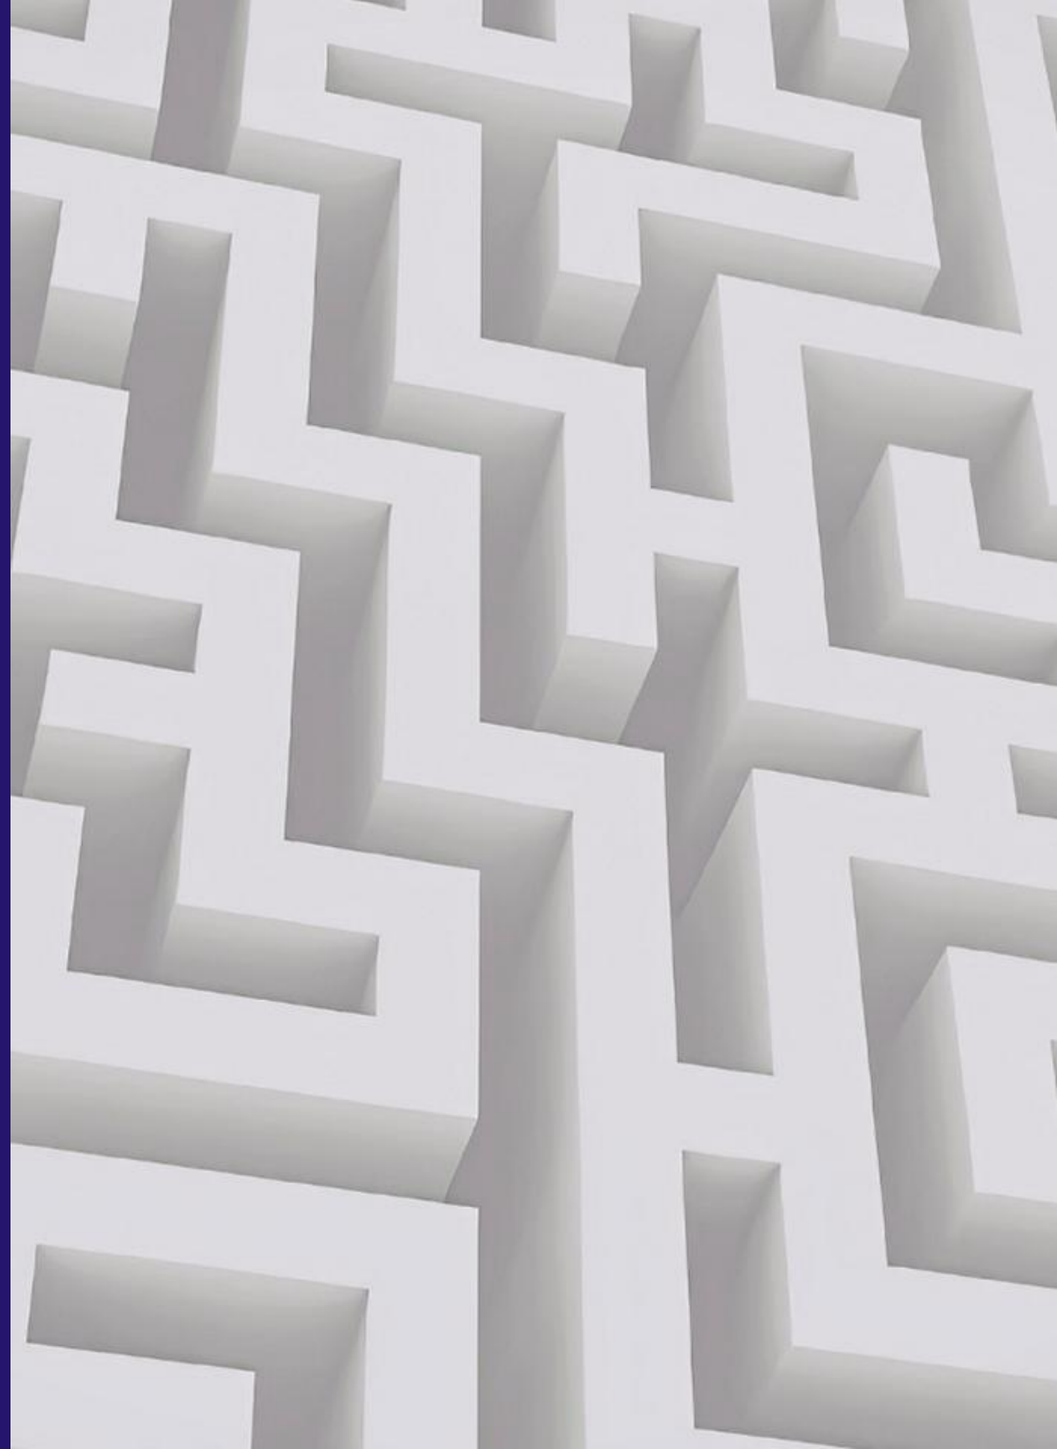


JOHN BUSH
FOUNDER, REVENUE EXPERTS AI


Our AI Strategy
AI-powered to be agentic
AI buzzword
autonomous and
another AI buzzword
yet another AI buzzword
we will accelerate
synonym for transformation
delivering
intangible outcome
at
really big adjective
scale.
Can we say "agentic" again?
We can never say it too much.
TOM FISH BURNE

Most companies already have an unofficial RAG strategy. It is just fragmented, inconsistent, and completely unmanaged.

So, if you are lost in a business AI maze, **context is your exit.”**

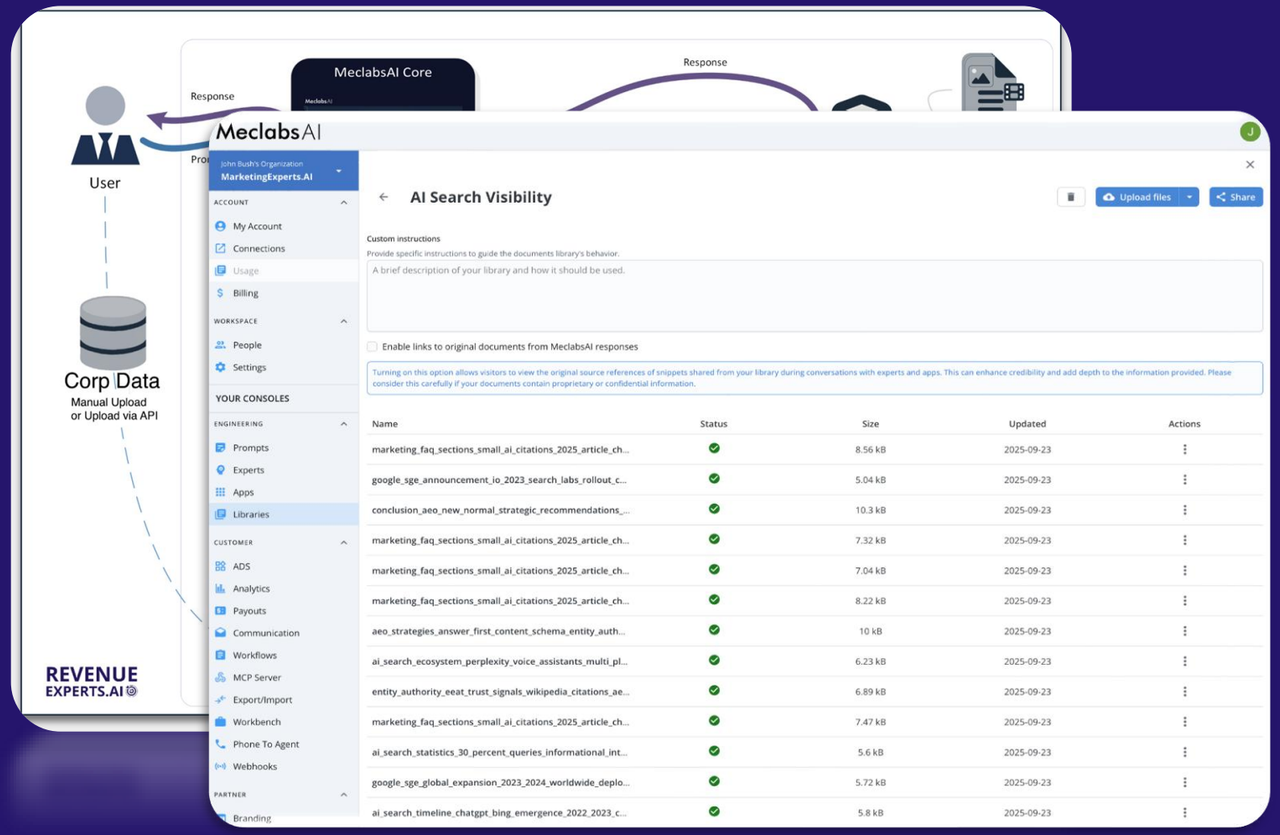


Three Terms to Ground our Conversation

- **Knowledge Base**
  - A prepared, structured, and governed collection of company information that AI can reliably find, interpret, and use.
- **RAG (Retrieval-Augmented Generation)**
  - A method where AI first retrieves relevant information from that knowledge base, then uses it to generate a grounded answer
- **Context Engineering**
  - Organizing, structuring, and shaping information so AI gets the right context, in the right format, at the right time.

From Experiment to Excellence

Our 2½ Year Journey - 3 things we have learned...



1. Keep your data safe
2. Organize and structure it for “Ingest” & “Retrieval”
3. ***Deliver better experiences with grounded answers***

Andrej Karpathy

@karpathy

"Something I'm finding very useful recently: using LLMs to build personal knowledge bases for various topics of research interest."

19.8MM VIEWS ON X & 5,000+ Stars



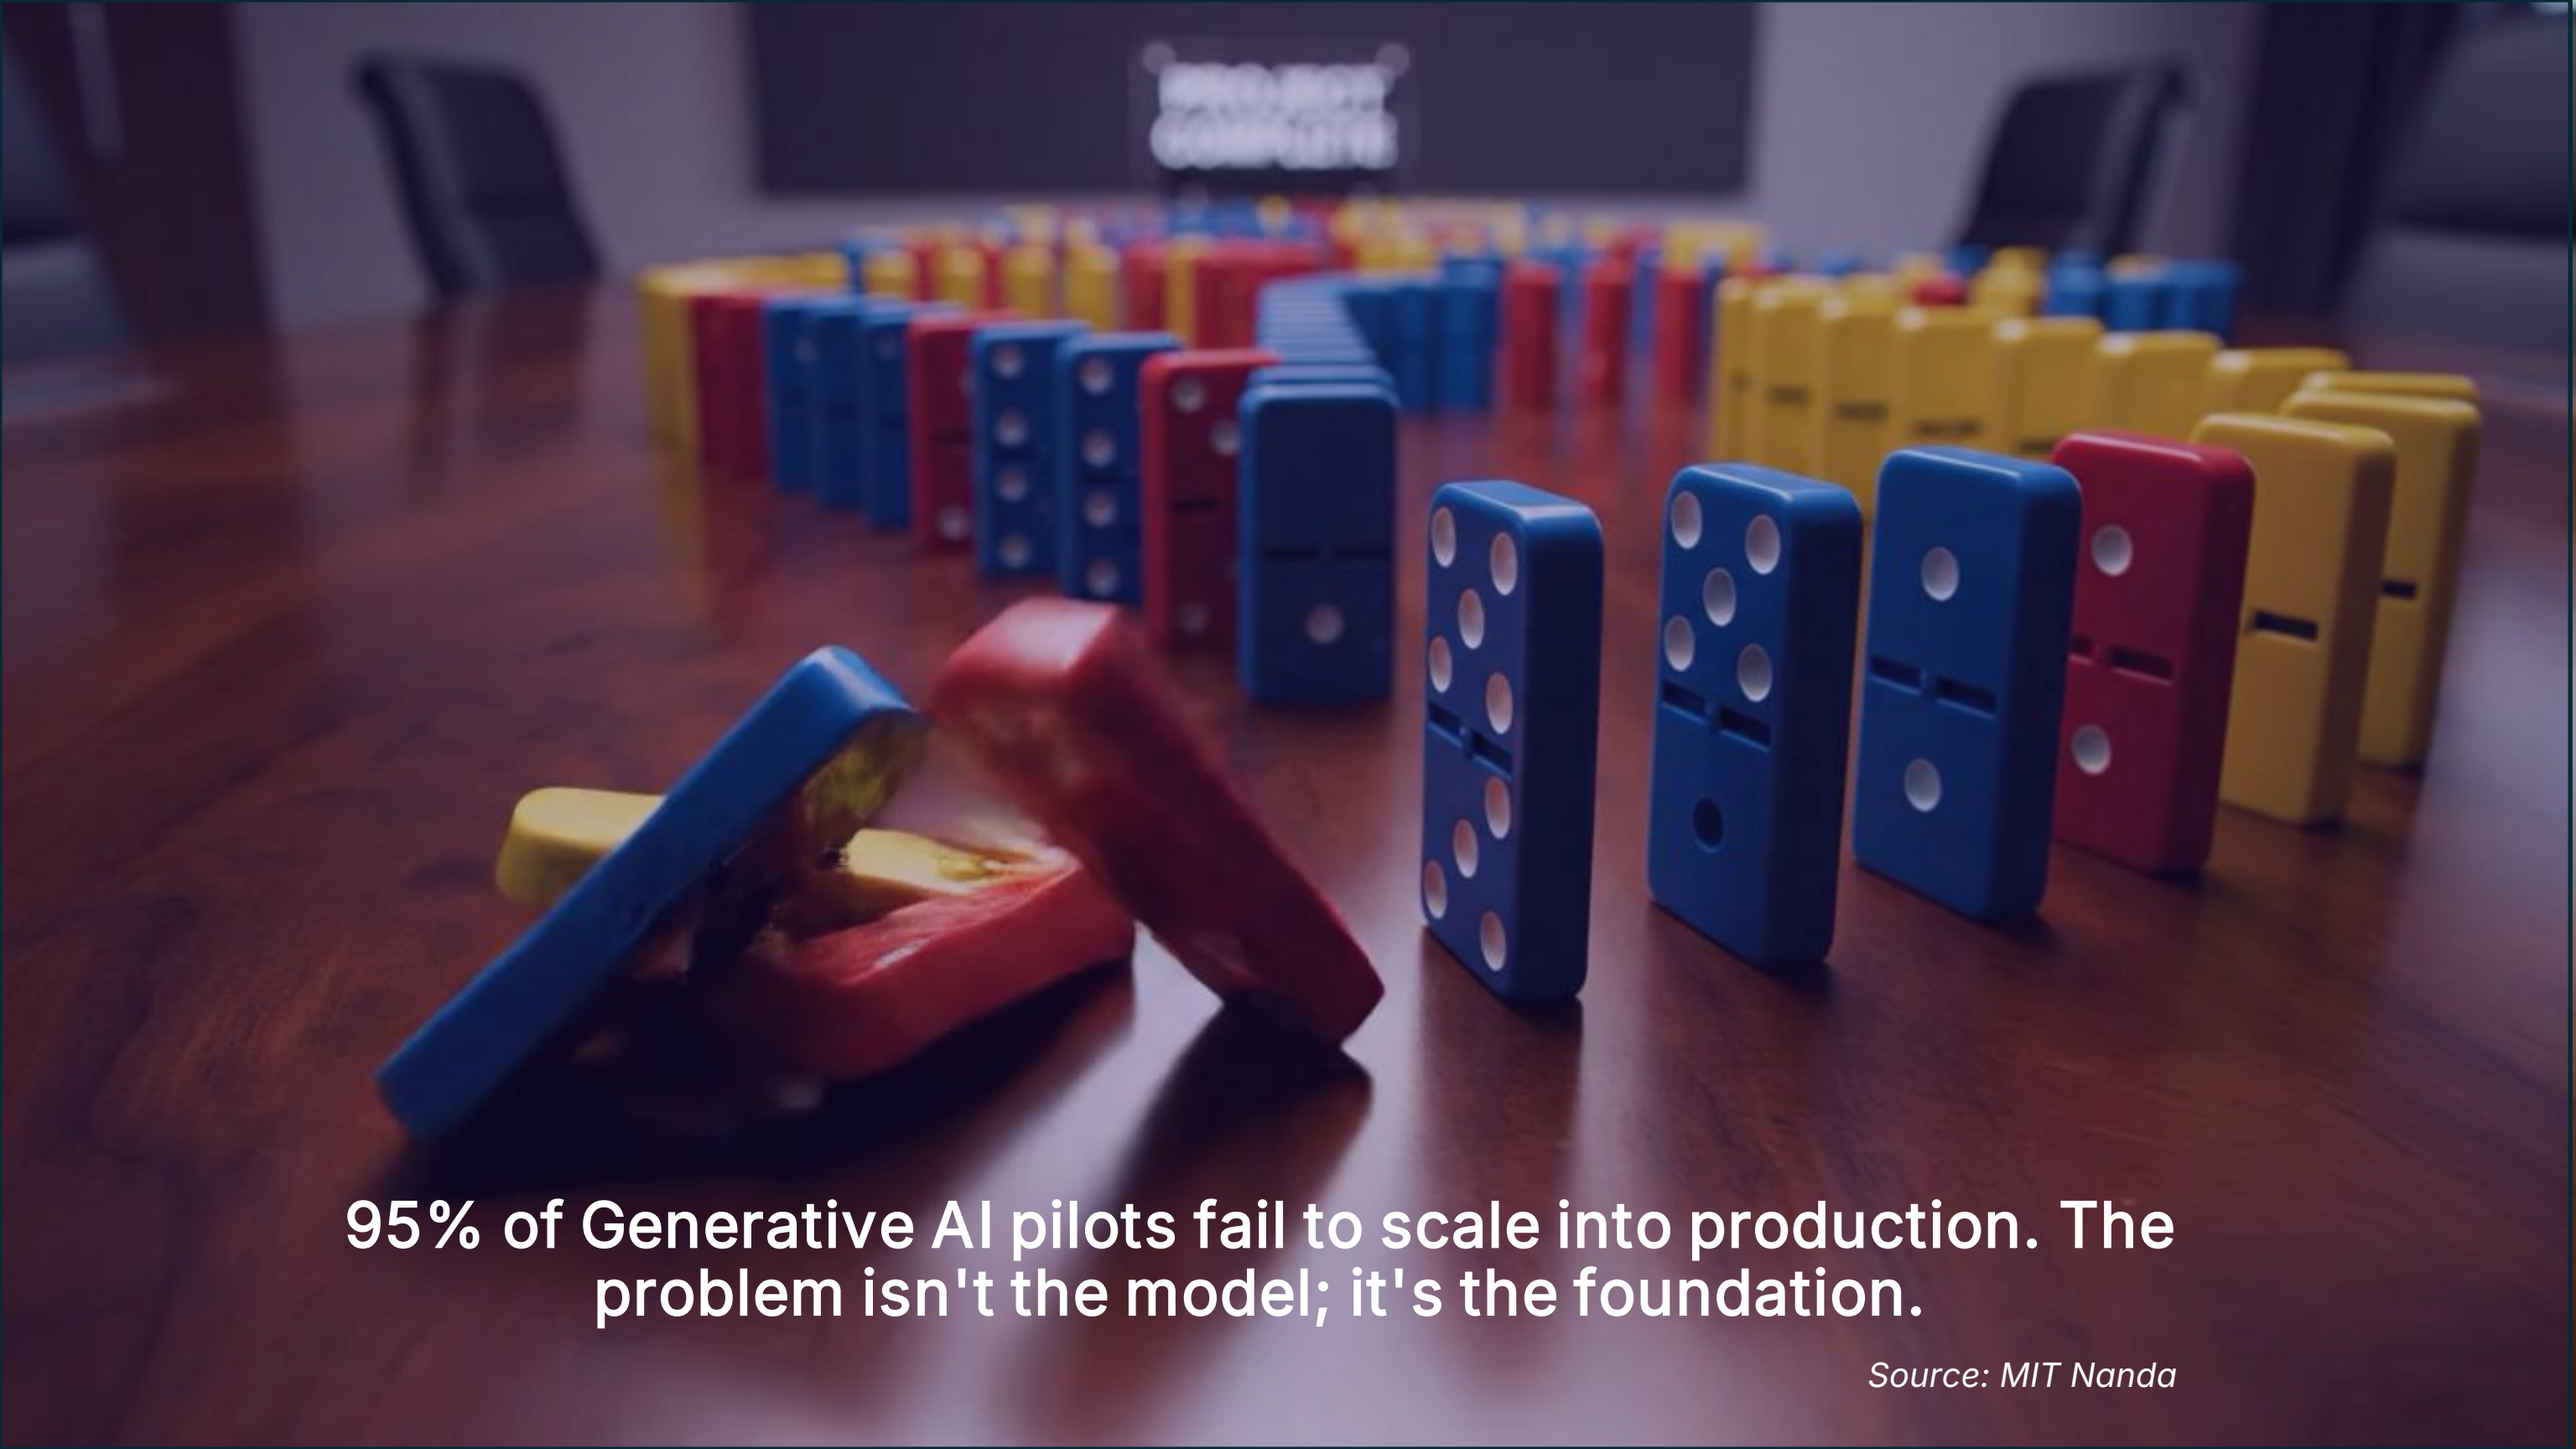
Automated Ingestion

LLMs autonomously compile diverse data into self-maintaining wikis, eliminating manual documentation and curation overhead

Strategic Intelligence

This framework enables complex, AI-driven analysis and Q&A, effectively turning your research into a database.

Enterprise Knowledge Compilation Layer



95% of Generative AI pilots fail to scale into production. The problem isn't the model; it's the foundation.

Source: MIT Nanda

It’s a Real Problem: And Business AI just feels broken.

Most enterprise GenAI failures have not been due to model capability - they’re due to missing AI-ready context: governed, retrievable, current knowledge aligned to workflows



Preparing a Knowledge Base is not that Different

“You hire the smartest person you’ve ever interviewed. Knows everything about marketing, sales, and technology. But it’s day one. They know nothing about your company, your clients, or what your CEO said on a call last Tuesday.

So what do you do? You onboard them.”

So, think of your Internal AI like a brilliant new hire.
Onboarding AI with Structured Knowledge Repositories

Deployment Models Change. Context Strategy Doesn't

Cloud | Sovereign | Hybrid = delivery models

Cloud

Best for ▮ speed & scale

Use when ▮ low sensitivity, high volume

What you get ▮ Fast to deploy, latest models, lower upfront costs

Risk ▮ sensitive data exposure, limited control over retrieval context

Air-Gapped

Best for ▮ control & compliance

Use when ▮ regulated industries, IP protection

What you get ▮ Data sovereignty, no external exposure, maximum governance

Risk ▮ higher cost, slower iteration

Hybrid

Best for ▮ Real World Ent. AI

Use when ▮ you need both speed & control

What you get ▮ public models for reasoning, private retrieval layers for context, sensitive data stays

Outcome ▮ accurate, grounded, production-ready AI

The model is interchangeable **The context layer is your competitive advantage**

Think About Building Revenue Knowledge Systems

Turn your data into decisions that close deals

01 - INPUT

Ingestion Processing

- Documents, CRM, marketing, ops data, OCR, chunking, enrichment
- Pipeline control (Context Studio)

"Don't just connect data - understand it."

02 - INFO HUB

Knowledge Compilation

- LLM-generated summaries
- Entity extraction + relationships
- Versioned + auditable

"Turn documents into usable knowledge."

03 - SEARCH

Retrieval + Reasoning

- Hybrid retrieval (lexical + vector)
- Reranking & Structured joins

"Don't guess. Ground every answer."

04 - MONETIZATION

Revenue Intelligence

- Opportunity ranking
- Pipeline insights
- AEO-ready outputs

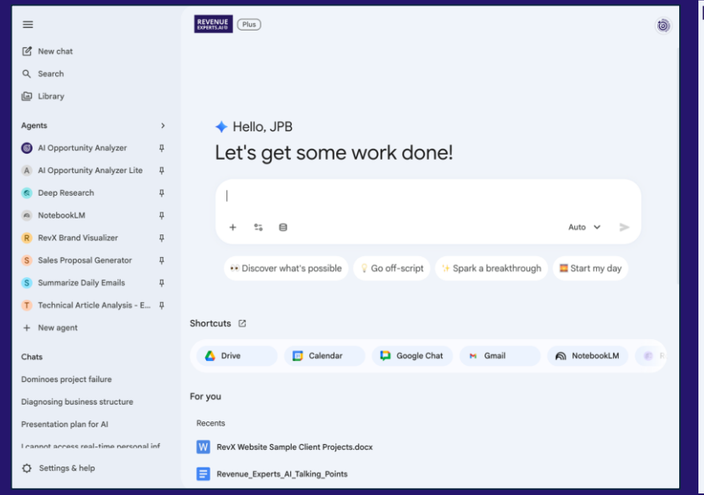
"This is where AI actually makes money."

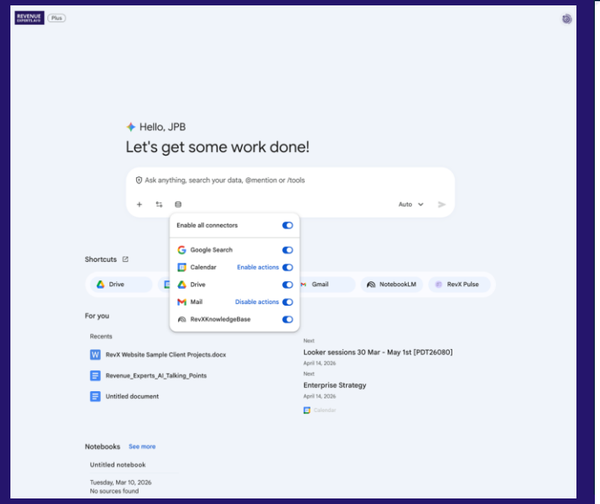
60% of AI projects will be abandoned through 2026 because they lack AI-ready, structured data, metadata + governance + observability + retrieval patterns

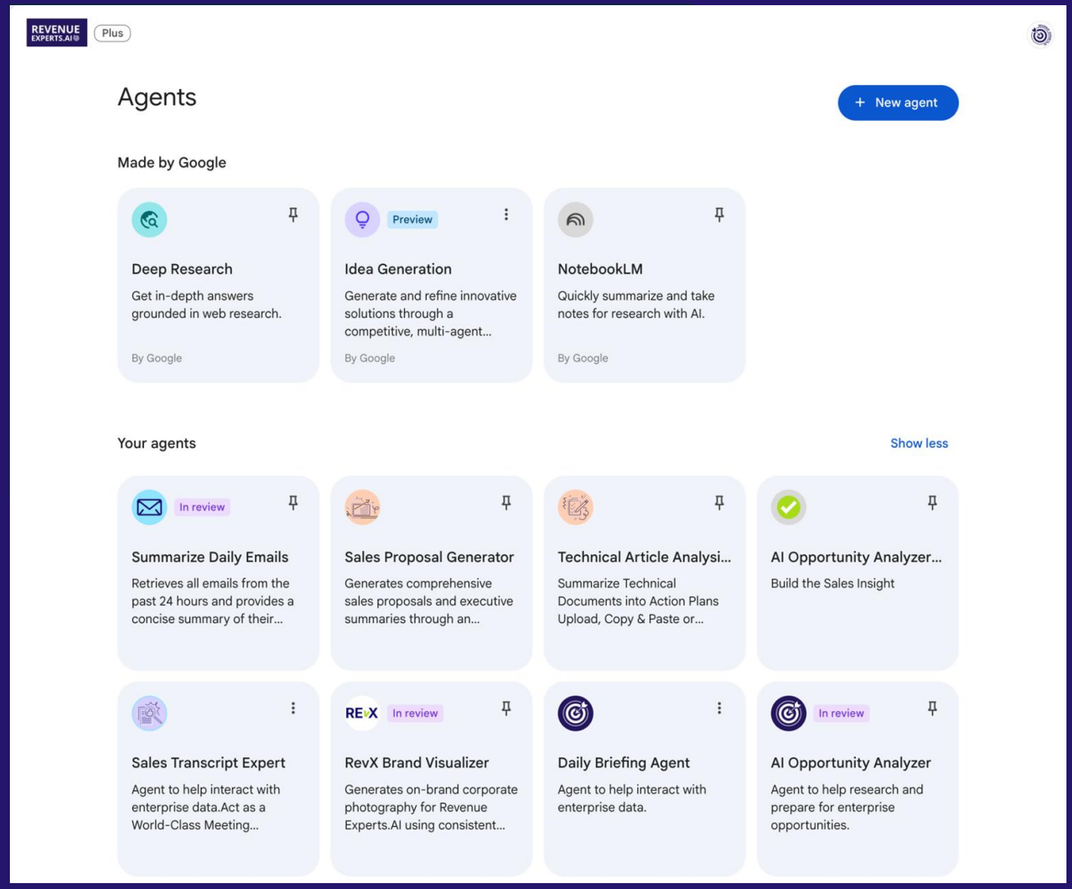
(Source: 2026 Gartner)

Gemini Enterprise - Agents Used Daily

Building automations & Workflows







Organized knowledge → Permissioned model → Workflow clarity → Grounded retrieval → Testing and oversight

Google Cloud

From Marketing Automation& Digital Selling to AI Context Engineers

How May You Ask?

Our Epiphany (Fall’24)

We needed to understand what was going on.

- We saw Google ad performance starting to soften
- Larger customers were questioning spend and channel effectiveness
- SEO was not dead, but buyer behavior was clearly changing
- People were moving from search results to AI-generated answers
- We realized AI rewards information that is clear, structured, and easy to retrieve

We needed evidence, not hype.
Our Vibe Coding was Born!

Using the right tools for the job

Methodology

We did not want to rely on tweets, demos, or AI hype cycles. Instead, we discovered LLMs actually accessed, interpreted, and cited the web.

36 Point Rubric

A rigorous framework to evaluate AI visibility, authority, structure, retrieval, and citation readiness-the factual foundation for our engine.

Enterprise Stack

T
Tested behavior using proven enterprise tools and controlled workflows

Toolchain

- Replit & Cursor (Dev)
- n8n & Make (Automation)
- ScraperAI & Firecrawl (Ingestion)

Humans and LLMs don’t read websites the same way.

Proof: The AI Search Visibility / AEO Issue

The Problem

Beyond traditional SEO, their brand was flying blind.

They had no idea if LLMs like ChatGPT or Perplexity were referencing them when buyers asked for recommendations.

The AEO Solution

Comprehensive AI Search Visibility Analysis.

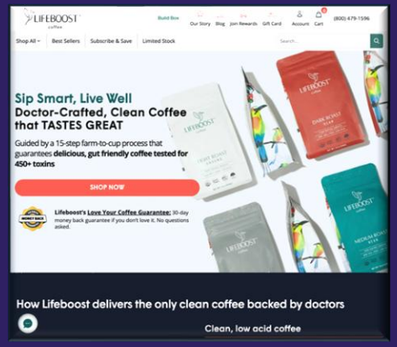
Baselined LLM presence and restructured content into a knowledge-first, AI-readable format.



The True ROI

2x Revenue Growth

Replaced guesswork with a data-driven AEO roadmap, exposing gaps and driving massive scale in 2026.

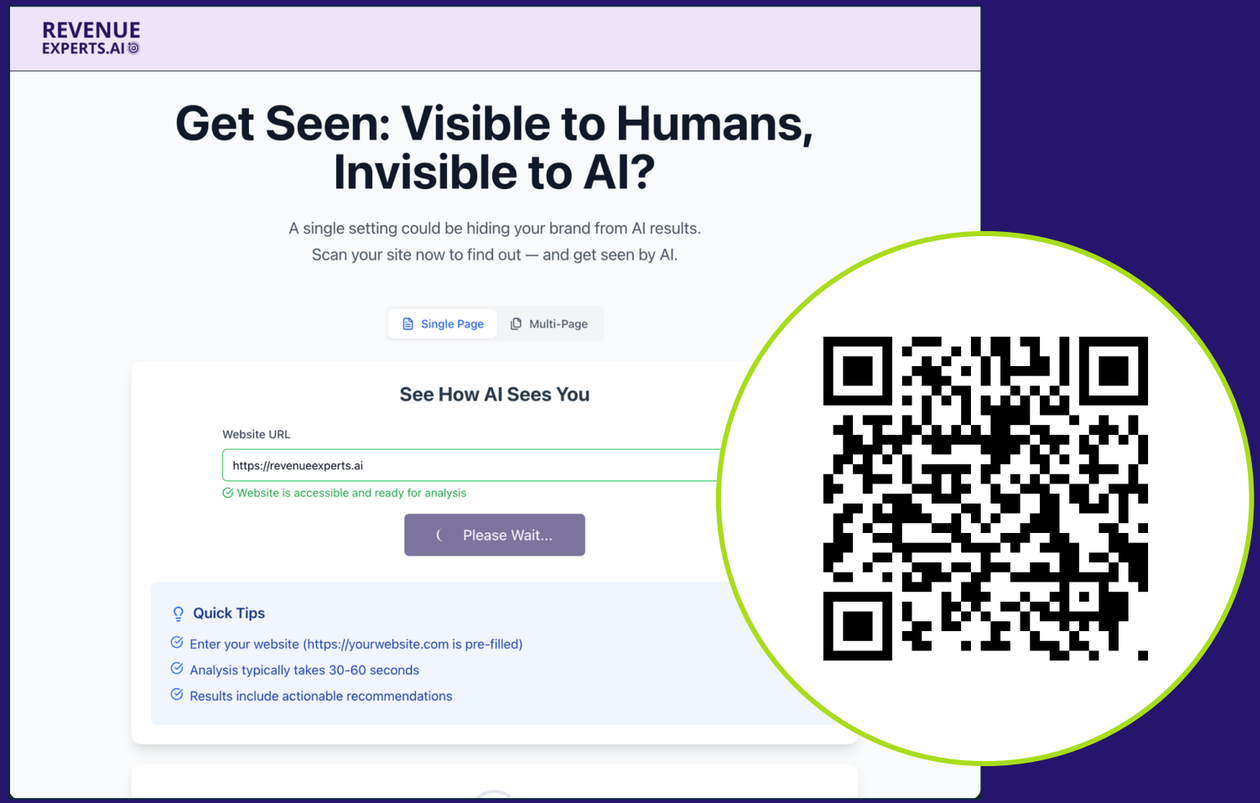


Case Example: AI Search Visibility results for lifeboostcoffee.com

In Principle: Y*ou increase the probability of being cited-if your content is authoritative, structured, and current."*

Actionable Insights from an Audit

- **Be Readable:** Identify specific content structures (Schema, Q&A, Technical NLP) to make your site AI-digestible.
- **Control the Narrative:** Prevent LLMs from hallucinating pricing or features by grounding AI answers in your data.
- **Competitive Intelligence:** See which competitors are stealing your "Answer Share" and how to leapfrog them.



Check your AI Readiness: https://revx.llmauditpro.com

Use Special Code: Digimarcon50

Breakthrough: Our Epiphany

Answer Engine Optimization was the lightbulb to structuring enterprise data.

Public AEO (Marketing)	Internal RAG Optimization	Why It Matters for AI Retrieval
Optimizing for Featured Snippets / FAQs	Structuring Internal Q&A Pairs	AI struggles to extract answers from dense, 50-page PDFs. We must reformat internal knowledge into explicit Question/Answer formats.
Using Clear H1/H2 Heading Structures	Applying Semantic Formatting to Documents	If a human can't skim a document to find a section, the AI's chunking algorithm will also get confused. Clear headers create distinct boundaries.
Implementing Schema Markup	Tagging Internal Documents with Metadata	We must tag internal files with metadata (Date, Author, Client, Project Status) so the AI doesn't pull outdated data sheets.
Writing Concise, Direct Answers	Eliminating "Corporate Fluff" in Internal Docs	**"Garbage in, garbage out."** Train teams to write project briefs with extreme clarity to remove ambiguity and prevent hallucinations.

Case Study 2: From PDF Chaos to Structured Context

Challenge

Modernize a legacy application buried under **25,000+ siloed PDFs** and manual document intake.

The goal was not just extraction - it was to turn unstructured files into **usable, searchable, governed context**

Technical Approach

Instead of a one-off custom build, we connected the platform to a **repeatable ingestion and context engine**.

That pipeline handled document prep, extraction, structuring, and mapping into **SQL + Vector search**.

Scalability

Because the pipeline was productized inside the **4-layer model**, it became more than a project.
It became a **reusable ingestion framework** for future enterprise use cases.

ROI & Impact

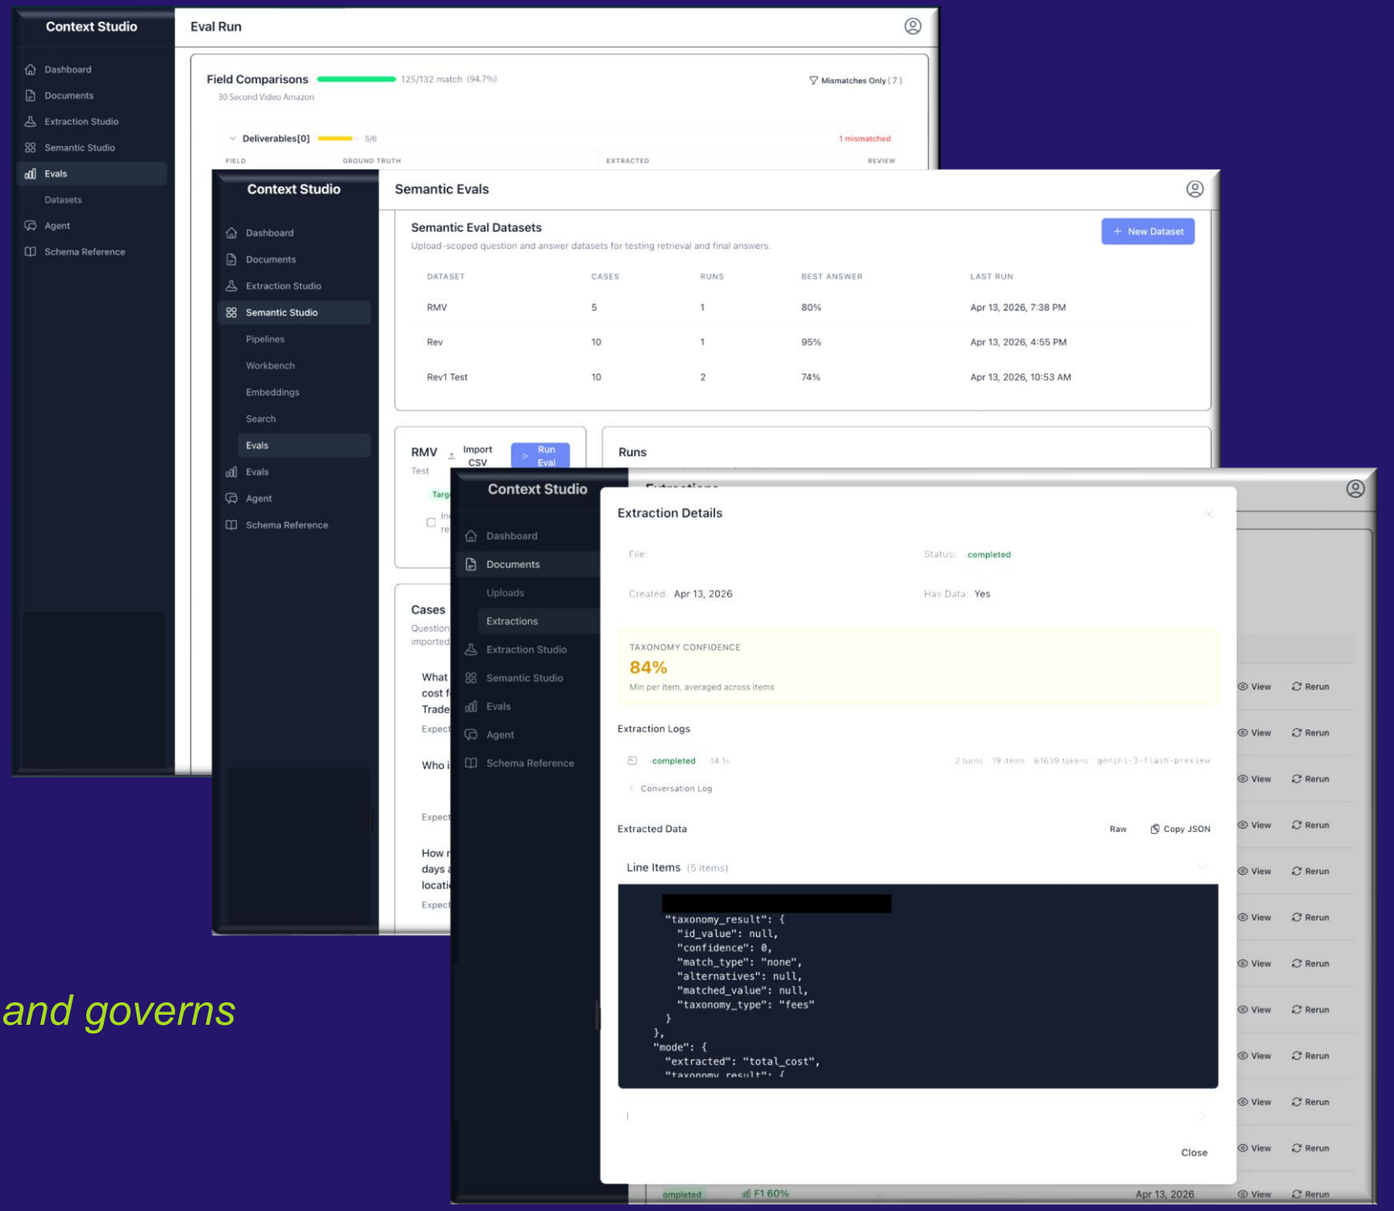
We removed the manual ingestion bottleneck, pushed extraction accuracy above **85%**, and achieved **80%+ semantic mapping**.

Human review shifted from data entry to **audit and validation**.

AI Context Studio

Because enterprise AI is only as good as the context behind it

- **Prepare** messy files for AI use
- **Structure** knowledge for retrieval and reasoning
- **Evaluate** outputs against expected truth
- **Govern** quality, freshness, and trust
- **Scale** grounded AI into real business workflows



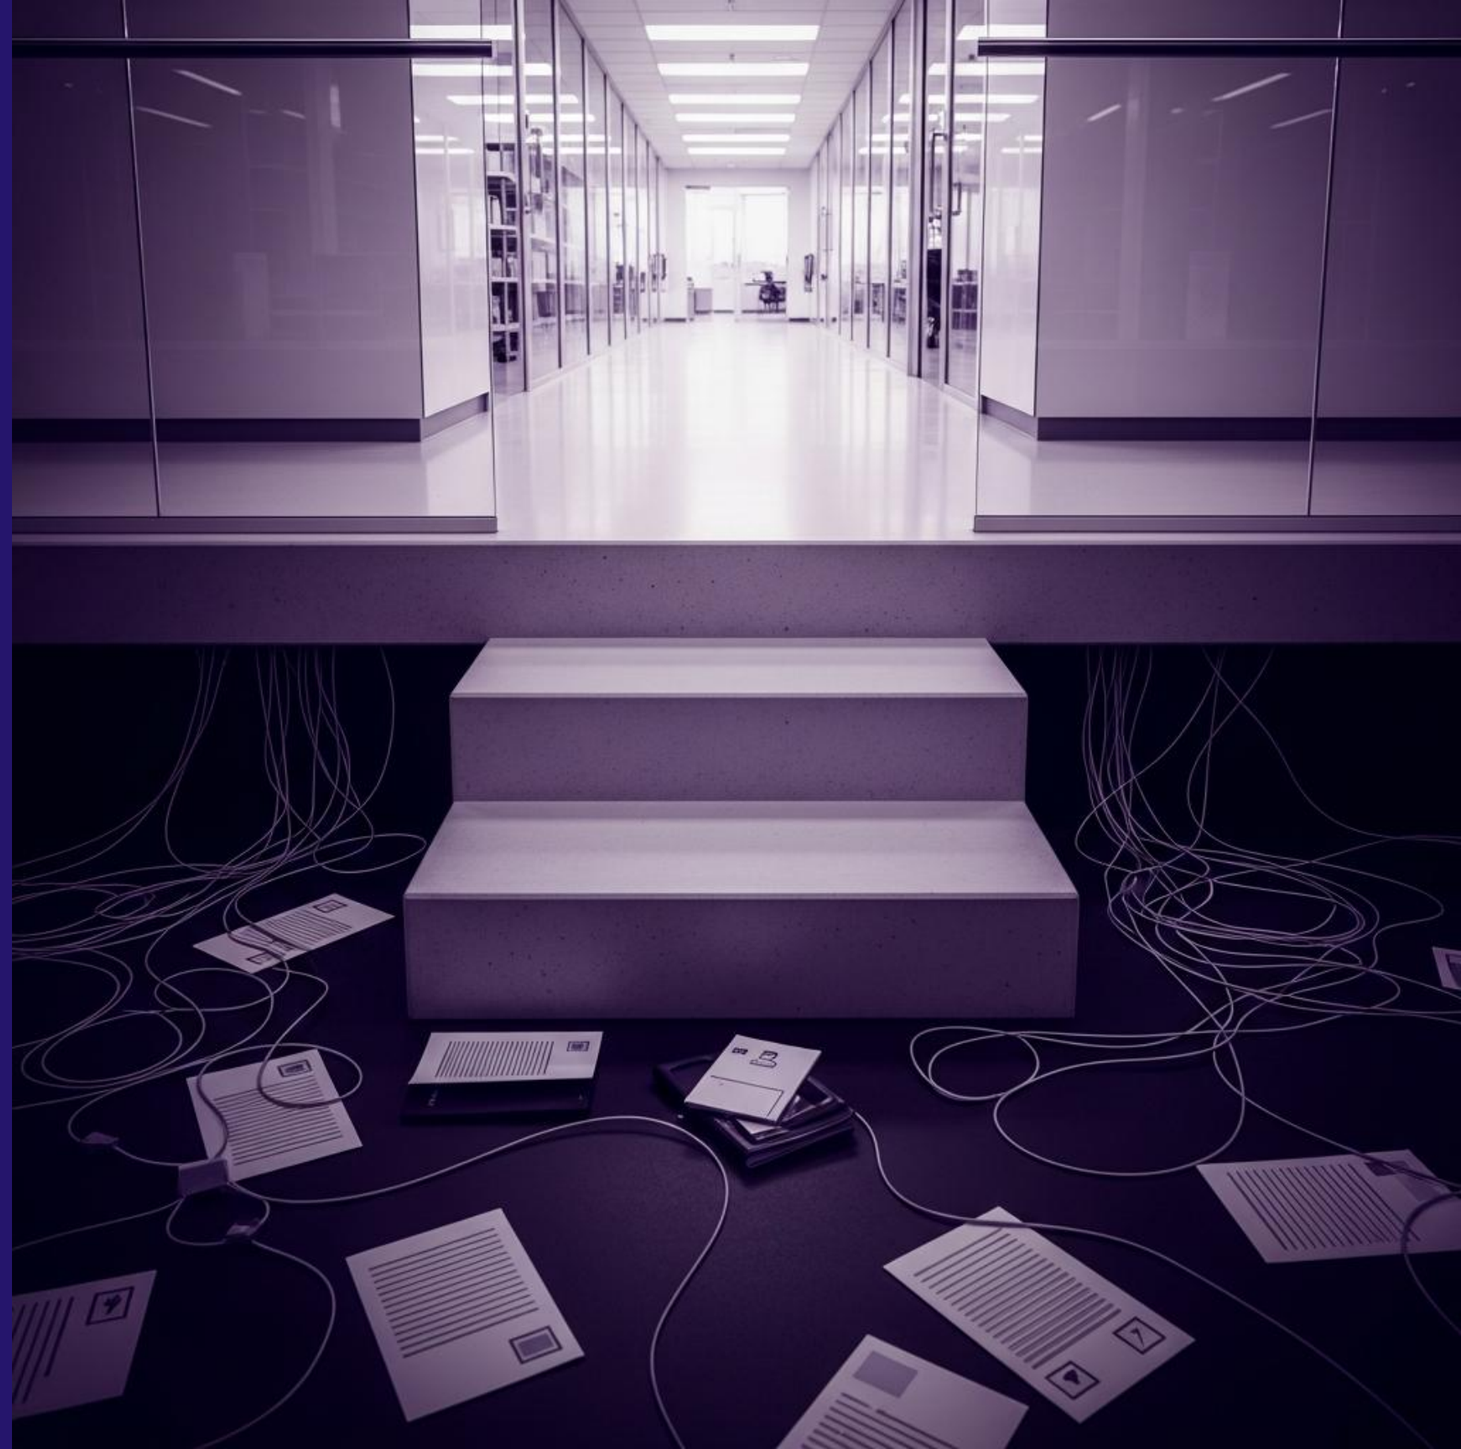
Context Studio is the layer that prepares, evaluates, and governs knowledge before AI uses it.

Google Cloud

Most financial firms are sitting on millions of PDFs with dense tables, charts, and exhibits containing information that is completely trapped… GenAI is changing this, but the problem is harder than the demos suggest.”

Source: Dean Barr, Head of Applied AI, The Carlyle Group 2026

If RAG is how AI looks things up, Context Studio is where you decide what is worth finding -
and whether it can enter the room and be trusted.



A Practical Framework

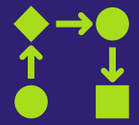
Three things to do. In this order.



01 Audit your content.



02 Structure it for retrieval.



03 Connect it to the workflows that matter.

Part 1 of 3 - Audit Your Content

Before you touch any AI tooling, map where your knowledge lives.

Content Type	Where It Lives	Last Updated	Owner	AI-Ready?
Product one-pagers	Google Drive / Shared	2025-Q4	Marketing	✗ No
Competitive battle cards	Slack pinned msgs	Unknown	Sales	✗ No
RFP response library	Email threads	2025-Q2	Sales Ops	✗ No
Customer FAQs	Help Desk + Drive	2025-Q1	Support	✓ Partial
Pricing & packaging	Confluence + PPT	Last updated: ?	Finance	✗ No

Most companies have never mapped this. You cannot connect AI to content you can't find.

Part 2 of 3: Structure It For Retrieval

Rewrite for Action

Rewrite key documents to answer specific questions.

Use FAQ Formats

Use FAQ formats for competitive, product, and objection content.

Create a Source of Truth

Build a repository for your most common sales/support questions. Start with 20.

Tag with Metadata

Tag content with metadata (product, persona, stage, date).

The same practices that make content more machine-readable internally also improve AI visibility externally.

Connect it to the Workflows That Matter

Do not deploy AI where it is possible. Deploy it where context changes the outcome.

Awareness - Marketing Execution

campaign briefs, persona research, content repurposing.

Consideration - Sales Enablement

pre-call prep, RFP responses, competitive objection handling.

Decision - Customer Engagement

support answers, onboarding guides, FAQ deflection.

Case Study 3: AgPal Proves the Point

The Challenge

Canadian agricultural SMEs faced a maze of fragmented federal and provincial funding information.

The Shift

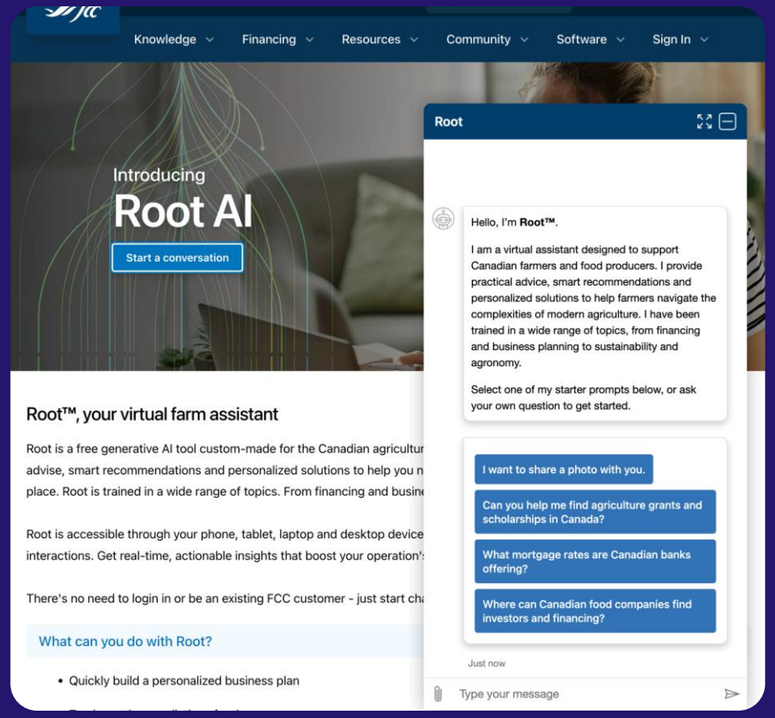
AgPal Chat used retrieval-augmented generation to ground answers in verified, curated program data.

Why It Worked.

The breakthrough came from data governance and architecture, not just the model.

The Real ROI

A trusted, practical B2B assistant that reduced hallucinations and made complex information usable.



So far, Root has supported 2,900+ conversations, with 91% of users reporting helpful results - plus nearly 400 producer-submitted images analyzed for equipment troubleshooting.

Source: StatsCan & Agriculture and Agri-Food Canada
Tool: agpal.ca

The Roadblocks to Scale

Three patterns we see repeatedly

The Chatbot Illusion
(Fragmented Knowledge)

Looks like progress. Hides fragmented knowledge.

- Buying the shiny AI tool and layering it over scattered shared folders, stale docs, and unmanaged content

The Reality

- Without a structured content foundation, you get fast answers - but not reliable ones

Your team spends more time checking answers than using them.

The Training Myth
(Ungrounded AI)

Trying to solve a context problem with model tuning.

- Believing you need expensive fine-tuning to get company-specific answers

The Reality

- Most teams do not need to retrain the model - they need to ground it in structured, governed knowledge

The system generates answers, but no one fully trusts them

The Business Disconnect
(Over Engineered Systems)

Built by IT, detached from the workflow.

- Large infrastructure projects built in isolation from the teams who actually needs the outputs

The Reality

- If it is not tied to a real workflow, a real owner, and a real business decision, it stalls

You built the system, but the business never adopted it

Overcoming Additional Challenges



Retrieval Misses

Wrong chunks retrieved during the process.

Mitigation

- Hybrid retrieval + reranking
- High-level non-technical test sets



Source Conflict

Stale or conflicting content to outdated answers.

Mitigation

- Source-of-truth rules
- Ownership & freshness SLAs



Prompt Injection

Malicious or hidden instructions inside retrieved content.

Mitigation

- Guardrails & least privilege
- Sanitize retrieved context



Overconfidence

Hallucinated answers with false certainty.

Mitigation

- Human-in-the-loop (HITL)
- Confidence thresholds

What this Unlocks

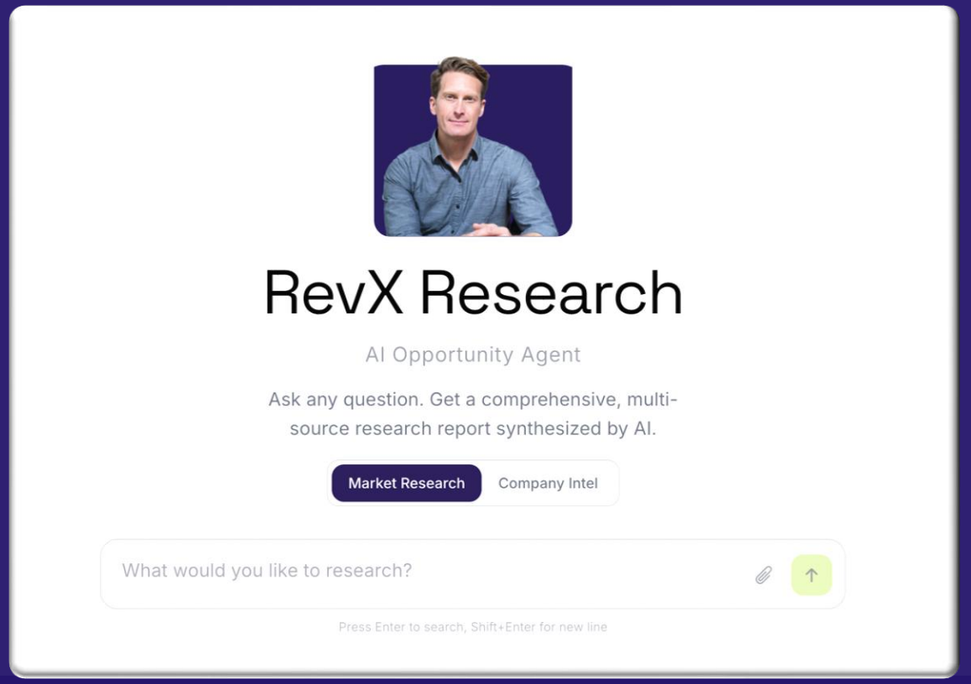


AoA: The Intelligent Sales Brief

Live search via Exa, grounded in our knowledge base, then analyzed with a reasoning model

From Search to Sales Insight

The AI Opportunity Agent is not a generic summarizer. It combines **live external retrieval**, a **curated RevX knowledge base**, and a **thinking model** to produce a grounded, decision-ready sales brief.



Explore the tool: www.revenueexperts.ai/aoa

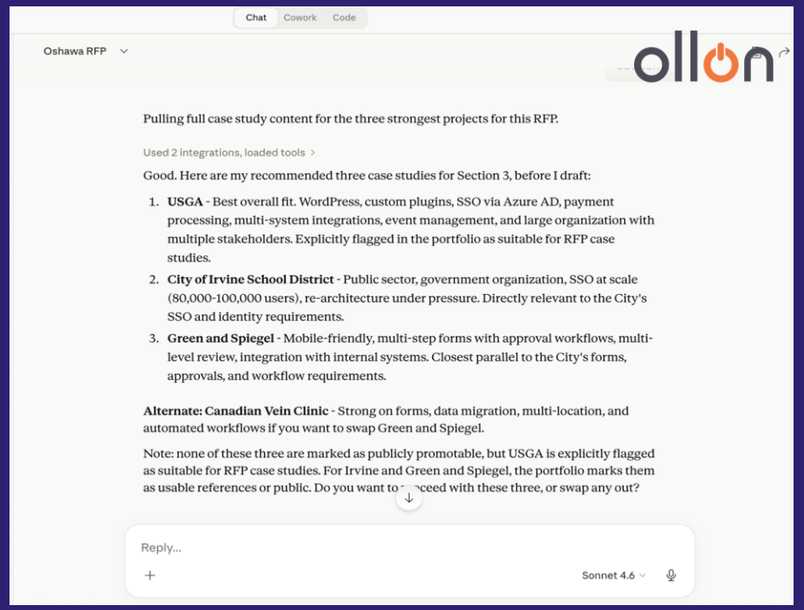
The Proposal Agent

Input Pipeline

- Call Transcript
- Client Wireframes
- Knowledge Base
- Process Docs
- Case Studies
- Document Guide

Finished Proposal

Output Results



- ✓ Consistent brand voice
- ✓ Grounded in trusted knowledge
- ✓ Tailored to the opportunity
- ✓ Client-specific talk track

2 hours

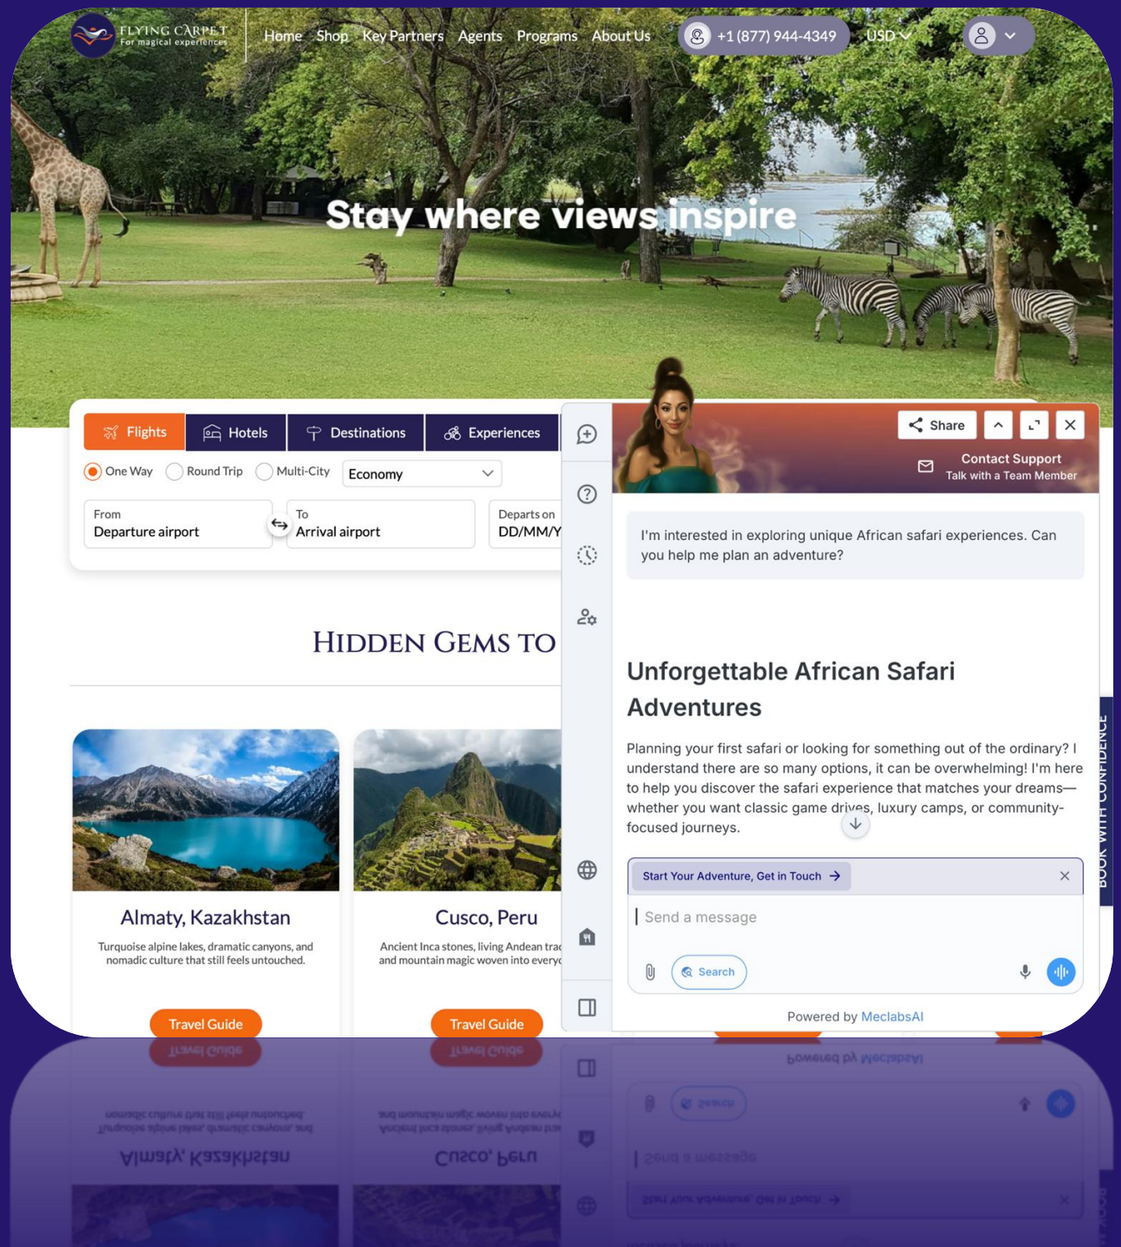
Conversational Agents

Context-aware AI that holds a conversation

Grounded, Guided and Auditable

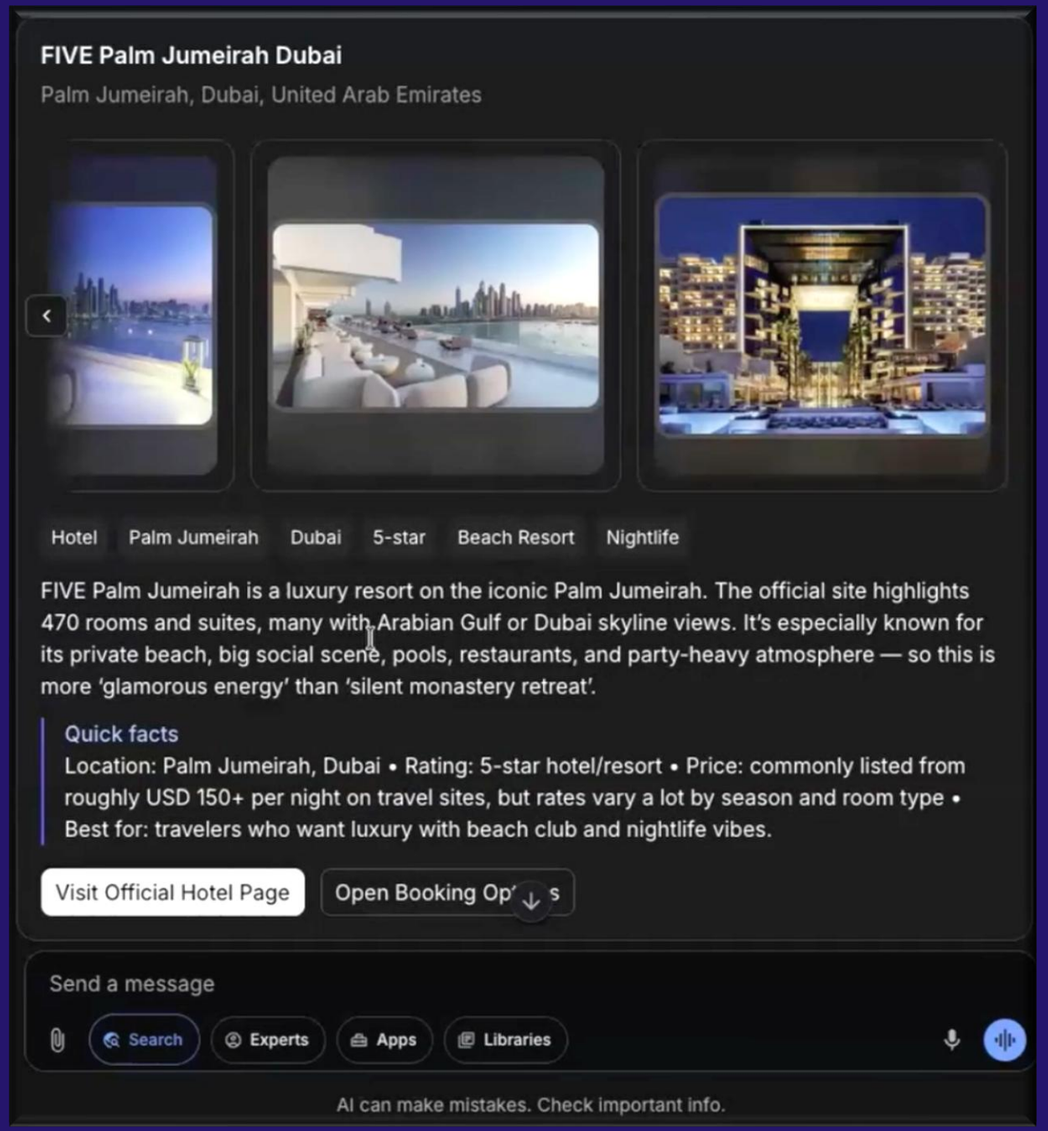
These agents are designed to support real business conversations by combining trusted knowledge, multi-turn memory, brand controls, and full auditability.

- **Grounded** in trusted knowledge
- Remembers context across turns
- Stays within brand and policy guardrails
- Logs responses and source behavior
- Generates analytics for improvement



Conversional Agents

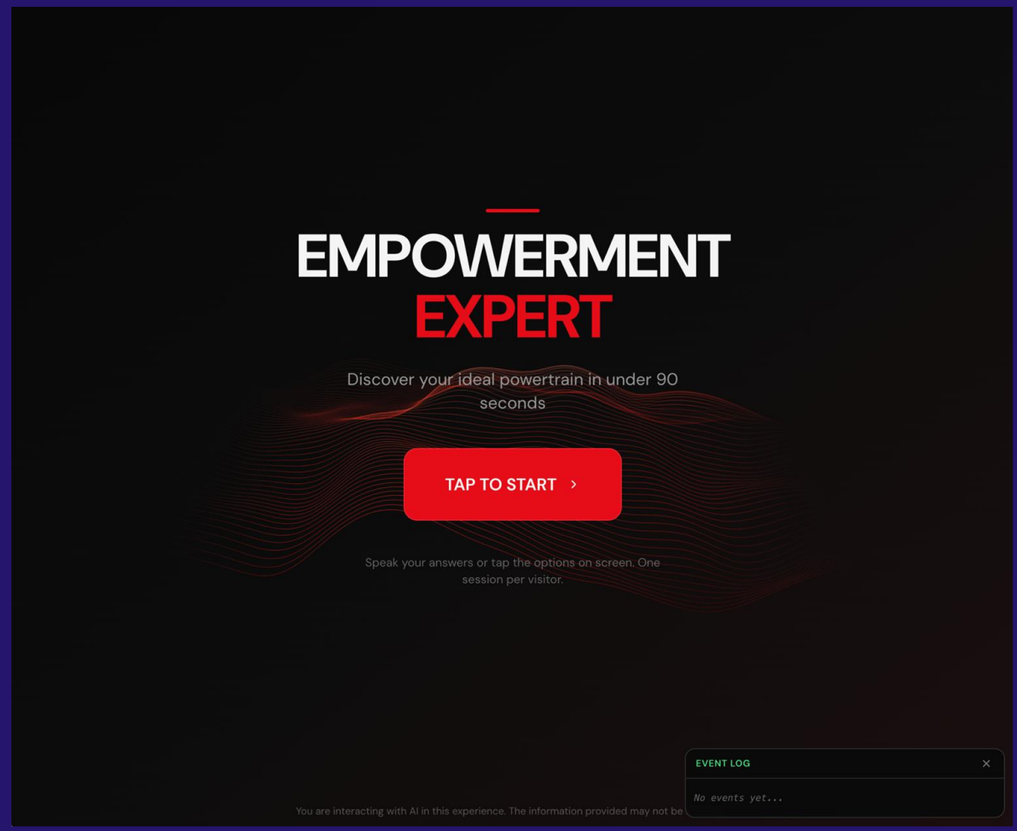
- Zero Learning Curve
- Experts, Apps & Libraries
- Built in Analytics
- Orchestrated Workflows



Brand-Safe AI in Production Environments

Built to operate inside real business rules - not just generate answers

- ✓ Constrained by brand, rules, and context
- ✓ Inputs isolated to reduce manipulation risk
- ✓ Outputs filtered before reaching the user
- ✓ Guided by defined and controlled workflows
- ✓ Every interaction monitored and auditable



Don’t Let Content Go to Die

1. The CRM

Full of notes no one formats consistently, attached PDFs, and stale contact records.

2. The Shared Drive

4,000 files, no naming convention, and nobody knows what's current.

3. Employees' Heads

Institutional knowledge trapped in Sarah's Slack messages and Tom's email drafts.

The High Cost of Information Chaos

The Internal AI Partnership

IT Owns the Engine, the Business Own the Fuel

The AI Myth

"AI is an infrastructure problem."

Reality: IT owns the pipes, but Business owns the Context - content, messaging, customer data.

The 95% Trap

Projects fail because they are built without structured context. The model is smart, but it's empty.

Your Opportunity

You don't need a massive IT overhaul to start.

- Organize your information today
- Launch targeted pilots
- Become the AI thought leader for your enterprise

Frame your data as the AI's source of truth

Start by organizing and structuring what the business already knows.

What to Do Tomorrow Morning

Three Concrete Actions, No new software required



1. Map Where Knowledge Lives

List your top 20 documents:

- Where they are
- Who owns them
- Last updated / AI-readable

½ Day



2. Pick One Use Case

Don't boil the ocean:

- Sales objection handling
- Support FAQ deflection
- Campaign brief generation

Pick one workflow, do it well, build the case for the rest.

One Decision



3. Readability Audit

Paste top 5 product pages into ChatGPT / Gemini / Claude and ask:

"What do you know about us?"

If the answer is wrong, thin, or outdated - fix content before strategy.

20 Minutes

Q&A

One question I hear a lot:

"How do we know if we are actually ready for AI?"

John Bush
john@revenueexperts.ai · revenueexperts.ai

Back Up

Agents

- Way too early for the Enterprise
- Numerous security and unmanaged issues

Author: Brian Binovsky

As Nvidia has decided to lay down their standard for OpenClaw in the Enterprise let's focus on that instead of OpenClaw, IronClaw or NanoClaw.

Stability and security problems with NemoClaw:

1. It has no controls to prevent DNS level directs.
2. It allows the user to control what external resources it can access, but does nothing to filter the responses from those sources. This easily leads to the risk of prompt injections. If the target resource is flagged as malicious the restriction is not reflected in the access control at all. Sometimes the LLM will still go there unrestricted.
3. It does nothing about tainted skills in public skill repositories.
4. It does nothing about managing credentials the agents may have. It does nothing about the risk of an agent using other people's credentials in a multi-user environment (which OpenClaw was not really designed for a single system).
5. It does nothing about potential malicious or defective code that may appear in the codebase as it changes. Never mind just these risks in code imported from external repositories on the Internet, it applies to code Claude generations as well.
6. It has no LLM token cost management of any kind. Not per user. Not for LLM tokens consumed by prompt injections.
7. Just because you put agents in containers, does not mean the containers are protected from all publicly disclosed security risks. So agents may be able to quickly exploit newly discovered risks - faster than humans can react to the threat.
8. It does nothing about the quality of repetitive tasks. Lacking any sort of project management construct, all actively is mostly bespoke. On 2 different days the process executed for the exact same request could be very different in more than the ways that


DIGIMARCON
CANADA 2026
DIGITAL MARKETING CONFERENCE