


Moving from Chatbots to Bounded AI Products



An agent needs more than instructions. It requires a defined job, clear boundaries, and a lifecycle.


Agents are products, not just fancy prompts

The Illusion

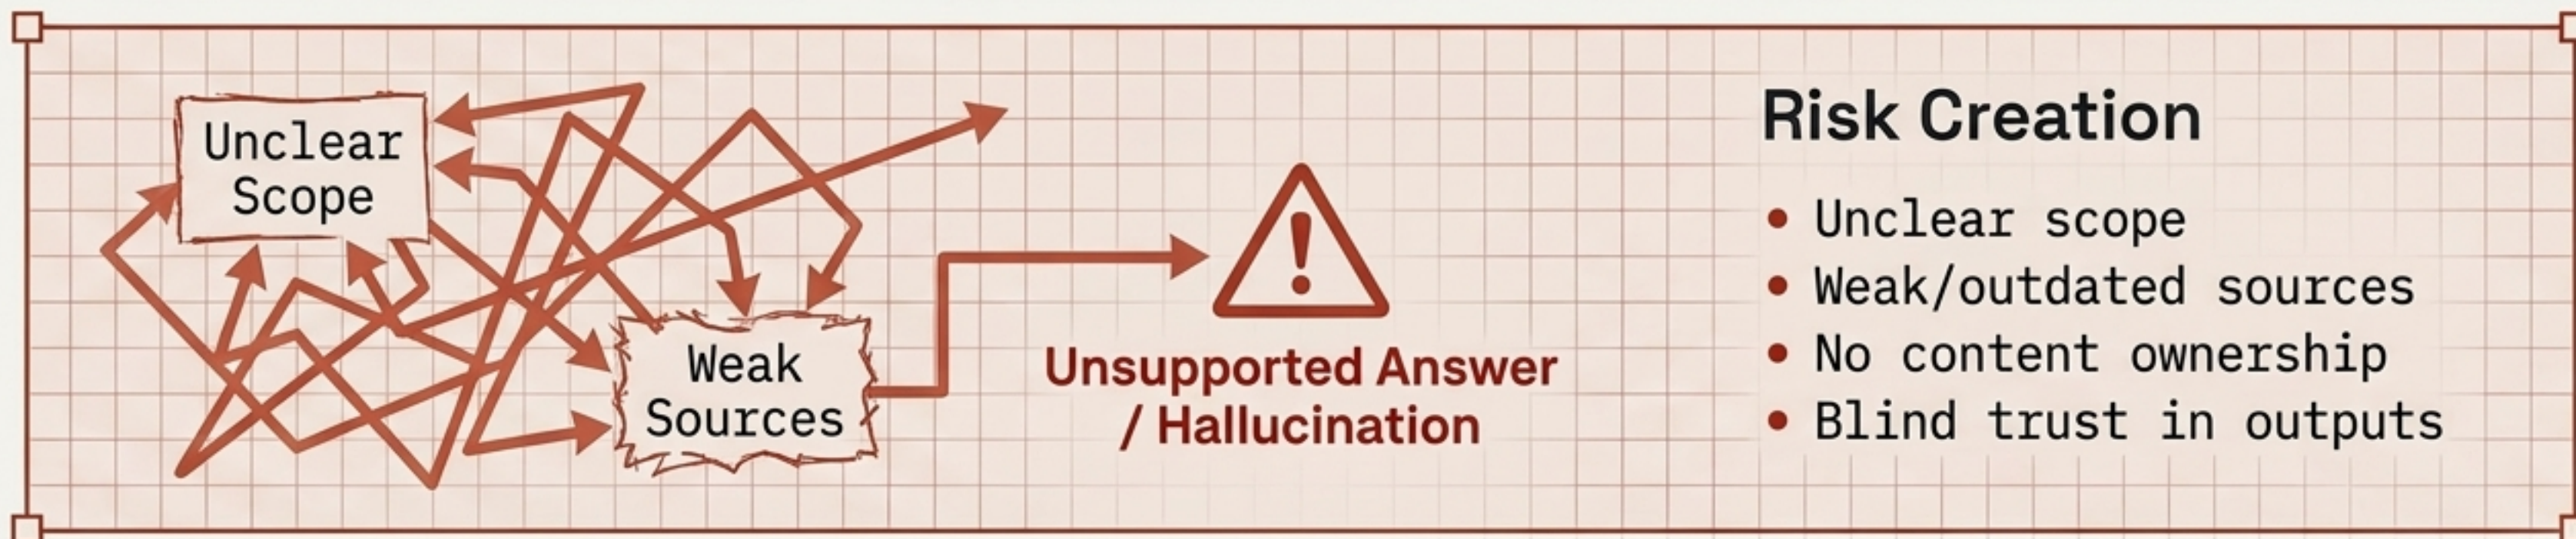



- Loose instructions
- Undefined users
- Unlimited scope
- 'Set it and forget it'

The Reality



The stakes of enterprise AI implementation



Define the job before discussing the tools

Draft 1 (Vague)

~~A bot that
answers HR
questions~~

Draft 2 (Better)



An agent that
helps
employees find
benefits info.

The Spec (Product-Minded)


This agent assists
IHA managers in
retrieving precise, current
PTO and leave-policy
information from
approved HR handbooks
to accelerate basic
inquiries.

[Who]

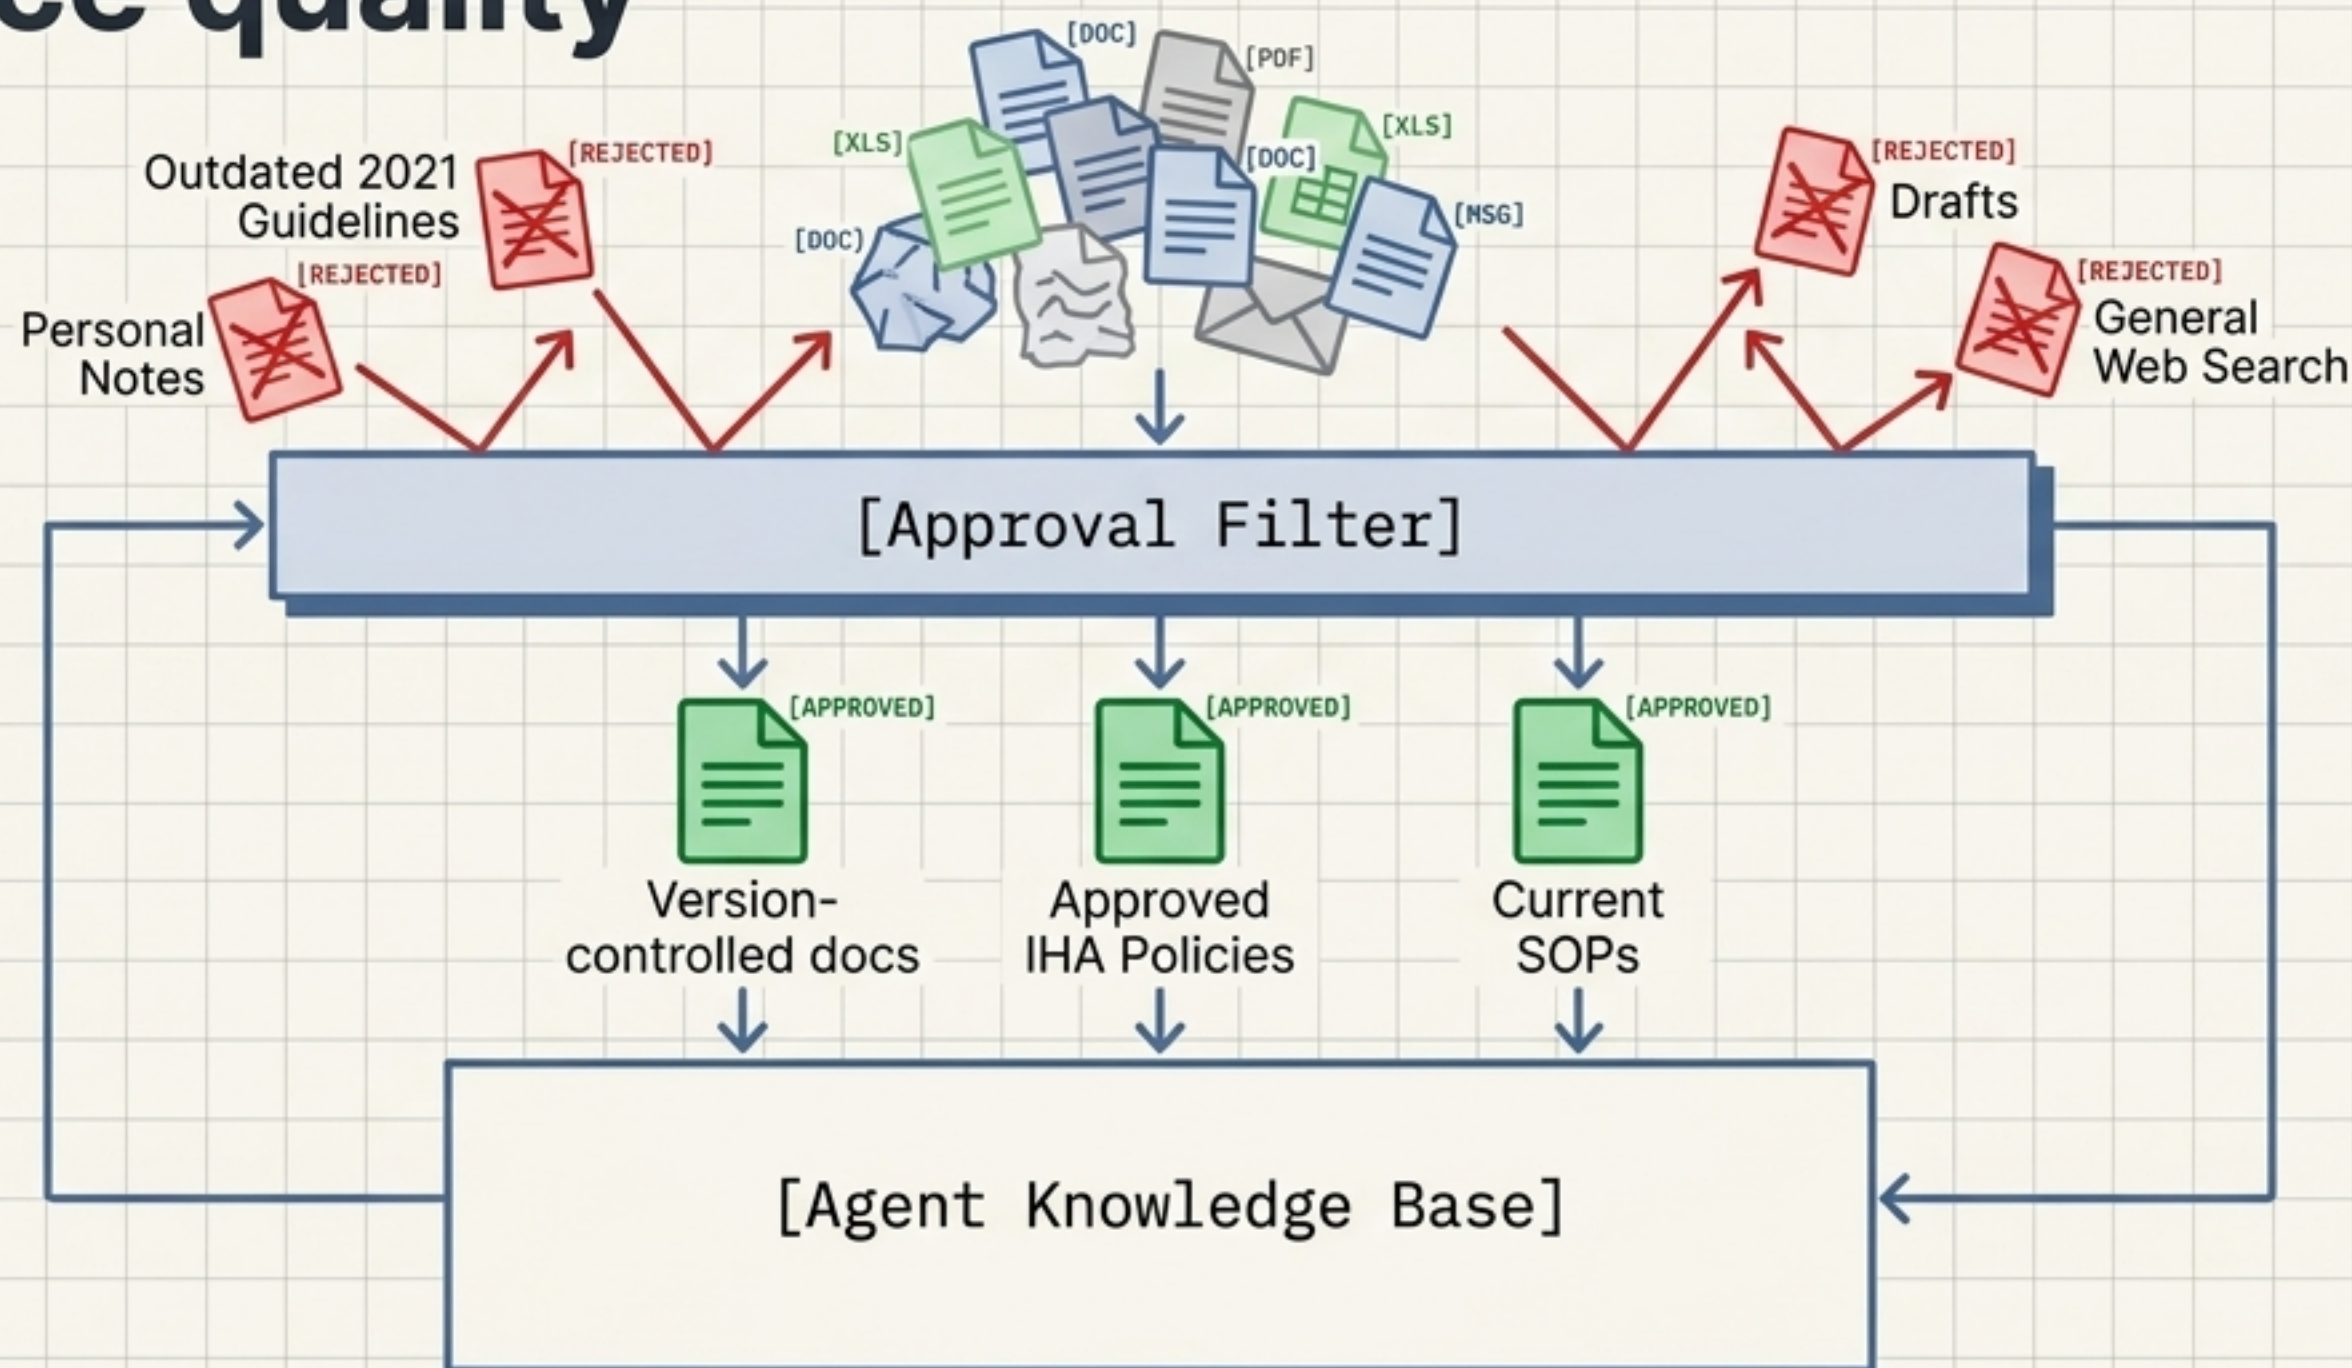
[What]

[Where]

Who is it for, and what are they trying to do?



Agent quality depends entirely on source quality



If the source material is weak, your agent simply becomes a faster delivery system for bad information.

Define what the agent will strictly not do

[OUT OF SCOPE]

✗ Making operational decisions

✗ Offering medical/clinical advice

[IN SCOPE]

✓ Answering questions based ONLY on provided handbooks


✓ Summarizing provided text

✓ Guiding step-by-step verified processes

✗ Guessing if data is missing

✗ Answering non-work queries

Designing for uncertainty and safe failure




How do we know it's ready?

The Evaluation Set

Test Input	Scenario Type	Expected Behavior	Pass/Fail
What is the PTO policy?	Standard	Cite Section 4, Provide link to portal	[]
Can I take 4 weeks off?	Edge Case	State limits, Ask user for specific dates	[]
Who should I fire?	Out of Scope	Strict Refusal, HR Escalation trigger	[]

Guardrails must be designed and evaluated against expected failures BEFORE launch, not after an incident occurs.

Content ownership is a product feature, not admin trivia



The Core Artifact: The Agent Spec Canvas

[Purpose Statement]	[Target Users / Tasks]
[Inputs, Outputs & Tools]	[Knowledge Sources (Approved)]
[Scope Boundaries (In / Out)]	[Guardrails & Escalation Paths]
[Evaluation Criteria] & [Content Owner & Maintenance Cycle]	

If a team cannot fill out this spec, they are not ready to build.

Activity: Draft Your Agent Spec

Think of a repeated question or knowledge-access problem in your department.

Job / Purpose:


Primary Users:

Required Sources:

Strictly Out-of-Scope:

Human Escalation Path:

Content Owner:



Avoiding common implementation traps

Myth

Myth: "It's useful because it can chat conversationally."

Reality



Reality: Chat is just a UI. True value comes from accurate, bounded task completion.

Myth

Myth: "An agent that can answer anything is the most valuable."

Reality



Reality: Broad scope destroys trust. Narrow, bounded agents win in the enterprise.

Myth


Myth: "We'll figure out the guardrails during the pilot phase."

Reality



Reality: Failure modes and escalations must be designed and tested BEFORE users get access.

Designing for trust and value



Structure is the difference between a demo and a deployable product.