



Achieving an 80% acceleration in time-to-value through an intelligent DDQ agent


Problem Statement

GPs require verbatim accuracy and native document fidelity when returning diligence questionnaires to LPs. The agent drafted answers quickly, but lacked export fidelity and audit workflows, forcing internal teams to manually clean up outputs before client delivery.

Role

Sr. Product Designer

Key Achievement: Compressed end-to-end turnaround time from 2-3 weeks to under 3 days



Ontra (DDQ Agent)

Challenge

A main challenge was resolving how legal and IR reviewers consume diligence data: whether they needed a **native document view** (maintaining spatial spreadsheet/Word layout) versus a **scannable list view**. In parallel, **engineering faced technical limitations** rendering dynamic text highlights and redlines directly inside document previews, requiring a lightweight interaction pattern to **verify source precedent without breaking review flow**.

Team

Sr. Product Designer (Myself), Sr. Product Manager, BE engineers, FE engineers

Scope/Constraints

To launch the MVP on a tight timeline, the product had to rely on existing **document viewer infrastructure** that couldn't support live redlines or bidirectional file editing. The primary **design consideration was to build immediate trust** in agent-generated answers despite these rendering limits, establishing **clear precedent attribution** while proving market demand before investing in a full custom editor.

Discovery

The Cost of Manual Diligence: Measuring the Pre-Agent Baseline

Collaboration: PM

2-3 weeks

Turnaround Time

Average calendar time required to return a completed 100-150 question diligence document back to an LP.

25-40 hours

Labor per questionnaire

Hours spent manually searching prior filings, copied spreadsheets, and legacy PDFs for first-pass drafts.

3+ teams

Stakeholder collaboration

Touchpoints spanning IT, Legal, Finance, and more to review, re-draft, and sign off.

Key Takeaway

The bottleneck in DDQ fulfillment wasn't typing the answers, it was the **operational friction of locating verified precedent** across fragmented systems and **managing asynchronous handoffs** between cross-functional teams.

Uncovering the Mental Models of Diligence Reviewers

Collaboration: PM

Building Confidence

- Reviewers need a quick way to validate the agent's findings
- Without one-click attribution they spend equal or more time fact-checking the agent

"If the agent populates an AUM figure or policy statement, I can't sign off on it until I see the exact LP report or side letter it came from. Fast drafts without clear sourcing give me anxiety."

Document Viewing

- Legal and IR directors insisted on seeing the holistic document layout
- Users evaluate how a response sounds relative to the question right before and after it

"Diligence questionnaires tell a story about our firm. We need to review answers in context of the full legal filing."

Providing Feedback

- Reviewers didn't want to rewrite the entire block of text manually
- Users wanted a quick way to instruct the agent to pull from a doc or adjust tone

"When the AI pulls an outdated clause, I shouldn't have to leave the flow or type a 400-word paragraph myself. Let me point it to our updated LPA and tell it to re-draft."

Translating Discovery into a Prioritized MVP Scope

Collaboration: PM, ENG

To bridge my research findings with engineering realities, I led a cross-functional alignment workshop to unpack workflow breakdowns and define a clear path forward.

- Shared Discovery & Breakdown Mapping:** Synthesized insights from user interviews to expose where the manual diligence process collapsed—specifically around silent hallucinations, opaque sourcing, and post-export cleanup.
- Affinity Mapping & Theme Clustering:** Grouped team observations into core recurring themes to build shared awareness around user trust, review speed, and precedent accuracy.
- Dot-Voting & Feasibility Matrix:** Facilitated team dot-voting across an Impact vs. Effort matrix to separate critical table-stakes from low-priority distractions.
- Aligning on the MVP Cut Line:** Established a strict Now / Next / Later framework, locking in source transparency and citation grounding as our required V1 while deliberately parking complex editing tools for fast-follow phases.

Common Themes

20 min. As a group, drag related stickies from Stage 3 up here and cluster them. Name each cluster with a theme header and merge duplicates. This is where scattered input becomes a short list of shared themes.





Solution

Transitioning from Legacy Pilot to a Scalable MVP

I leveraged Claude Code to stand up an interactive MVP workspace. This gave our team an artifact to test with stakeholders and validate directly with users.

- Navigation Architecture:** Replaced unclear, disconnected question tabs with a unified split-screen reading and triage layout.
- Attribution Transparency:** Transformed vague citation links into explicit rationale cards explaining *why* the agent selected specific precedent
- Cross-Functional Velocity:** Structured clear ownership, tagging, and approval states to remove bottlenecks during multi-stakeholder reviews.
- Active Feedback Loops:** Embedded lightweight correction and training triggers to continuously improve agent recall and accuracy over time.



RESULTS

The DDQ Agent delivered immediate impact following its MVP launch. Within the first 60 days, feature adoption surged **31% month-over-month**, while customers experienced a **near-80% reduction in turnaround time**, collapsing multi-week diligence cycles into just a few days.


Future State

A more intelligent and agentic partnership [View the prototype](#)

Building on workshop alignment and validation from our MVP launch, I designed our next evolution into an **interactive co-pilot workspace**. Grounded in concept testing, this design transitions the platform from a passive review queue into an active human-in-the-loop environment.

Key Design Enhancements

- **Tokenized Inline Citations:** Introduced numbered, sentence-level badges, linked directly to a source.
- **Direct Canvas Editing:** Empowered users to edit answers inline directly within the document preview.
- **Multilevel Agent Piloting:** Built a contextual feedback loop allowing reviewers to prompt revisions, attach precedent files, and co-pilot at both the question and document level.
- **Transparent Agent Reasoning:** Designed an explicit "Chain of Thought" activity log to provide real-time visibility into agent actions.



KEY LEARNINGS

Native Ecosystem Integrations

Meeting Reviewers in Their Existing Ecosystem

A key takeaway was that some enterprise legal and IR teams **resisted context switching** into a standalone web application for drafting. Rather than forcing users into a closed platform, high-velocity adoption requires **meeting reviewers where they already live**.

This insight sparked explorations for **native Word and Excel add-ins**, bringing Ontra's agentic precedent retrieval, citation matching, and verified library directly into enterprise desktop workflows.

Explainability & Reasoning

Elevating Agent Explainability and Audit Trails

While the MVP gave users a top-level summary of agent actions, high-stakes diligence **demand deeper provenance**. Reviewers needed granular visibility into *how* multi-source precedents were synthesized, parsed, and updated over time.

This directly motivated our future-state "Chain of Thought" activity log and chronological revision history, turning **ambiguous model behavior into an auditable, transparent decision trail**.