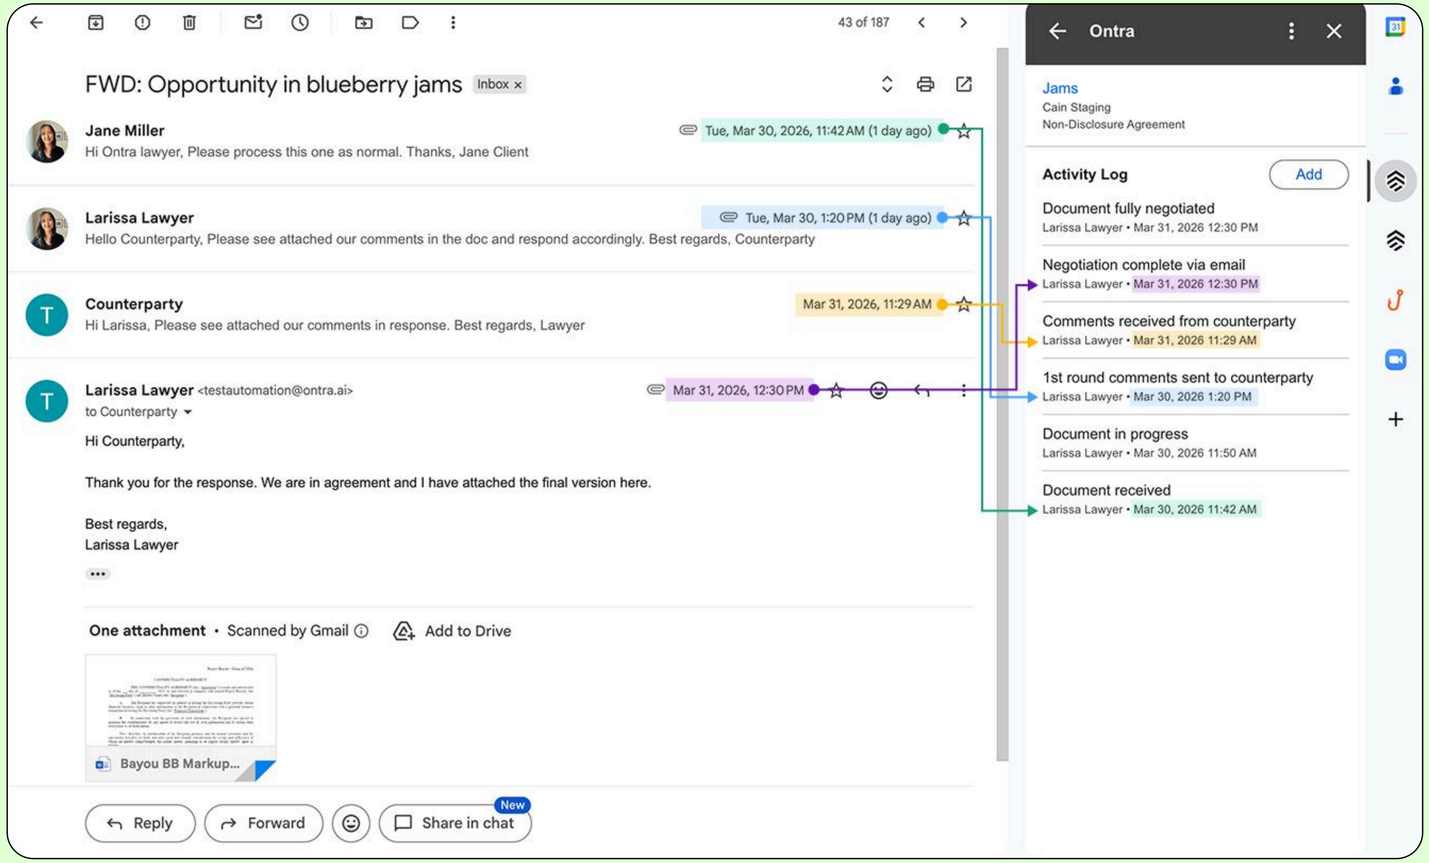
Capturing Real-Time Statuses from Emails with AI Classifications



Agenda

1. Context & Problem
2. Hypothesis
3. Research
4. Development
5. Results

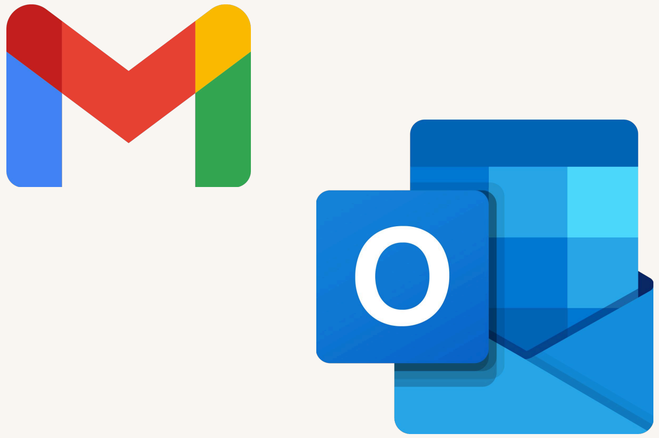
Company Background



Ontra is the AI platform for private markets – specifically for legal processes like negotiating NDAs in M&A deals.

- Founded in 2014
- Series B in 2021
- 1,000+ customers (private equity firms, investment banks, fund managers, etc.)
- 2 million+ contracts processed

The world of law and finance runs on email.

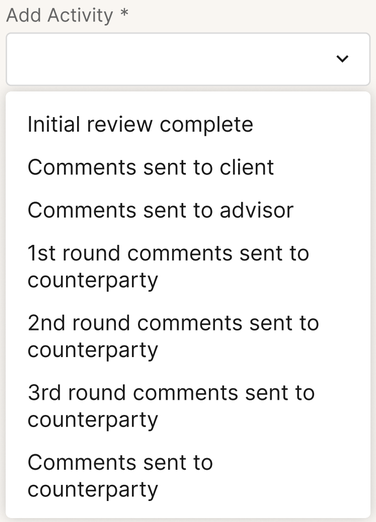


- Ontra's Contract Automation service runs on email because it fits into the existing workflows of our customers. Ontra's lawyers receive work requests, send comments to counterparties, and escalate to the customers — all in email.

- This works well, except it creates a giant blind spot for Ontra regarding the status of these documents. Ontra relies on mostly-optional, **manually-added statuses**, which means that in practice, most of the data is not added to the Ontra app.

- Q4 2025: Only 28% of documents had an email-verified tracked negotiation start timestamp, and 0% had first markup or fully negotiated timestamps.

Competition is heating up.



- Competitors were making bold claims — shorter turnaround times, lower prices — and winning on those promises. Customers were hearing the same pitch, and some were pushing back on Ontra directly, complaining about slow turnaround times and raising their expectations.

- The problem was that without accurate, complete timestamp data, Ontra had **no reliable way** to investigate those complaints, respond to them, or prove where performance actually stood.

HYPOTHESIS

We can get accurate timestamps for the negotiation lifecycle by going to Gmail.

- We knew the data lived in Gmail. We just needed a way to get it, as quickly and as accurately as possible.
- We decided on a hybrid approach of manual inputs and AI. We knew that the AI outputs would not be 100% accurate, so we wanted to build the muscle for users to complete manual inputs as well.

Dec 2025 – Jan 2026

Discovery

FINDING

In 7 research sessions, all users expressed the most positive sentiment toward the Gmail add-on option (vs. the in-app option).

FINDING

When users had a high volume of documents to negotiate, documenting statuses fell by the wayside.

QUOTE

“I like the Gmail plug-in option. Not having to go into the app would be great.”

DATA

80% of users use Gmail in the browser. (All users are given an Ontra Gmail email account.)

Hypothesis

This is a bet worth making.

DEVELOPMENT

Gmail Add-On v1: Manual Tracking

FUNCTIONALITY

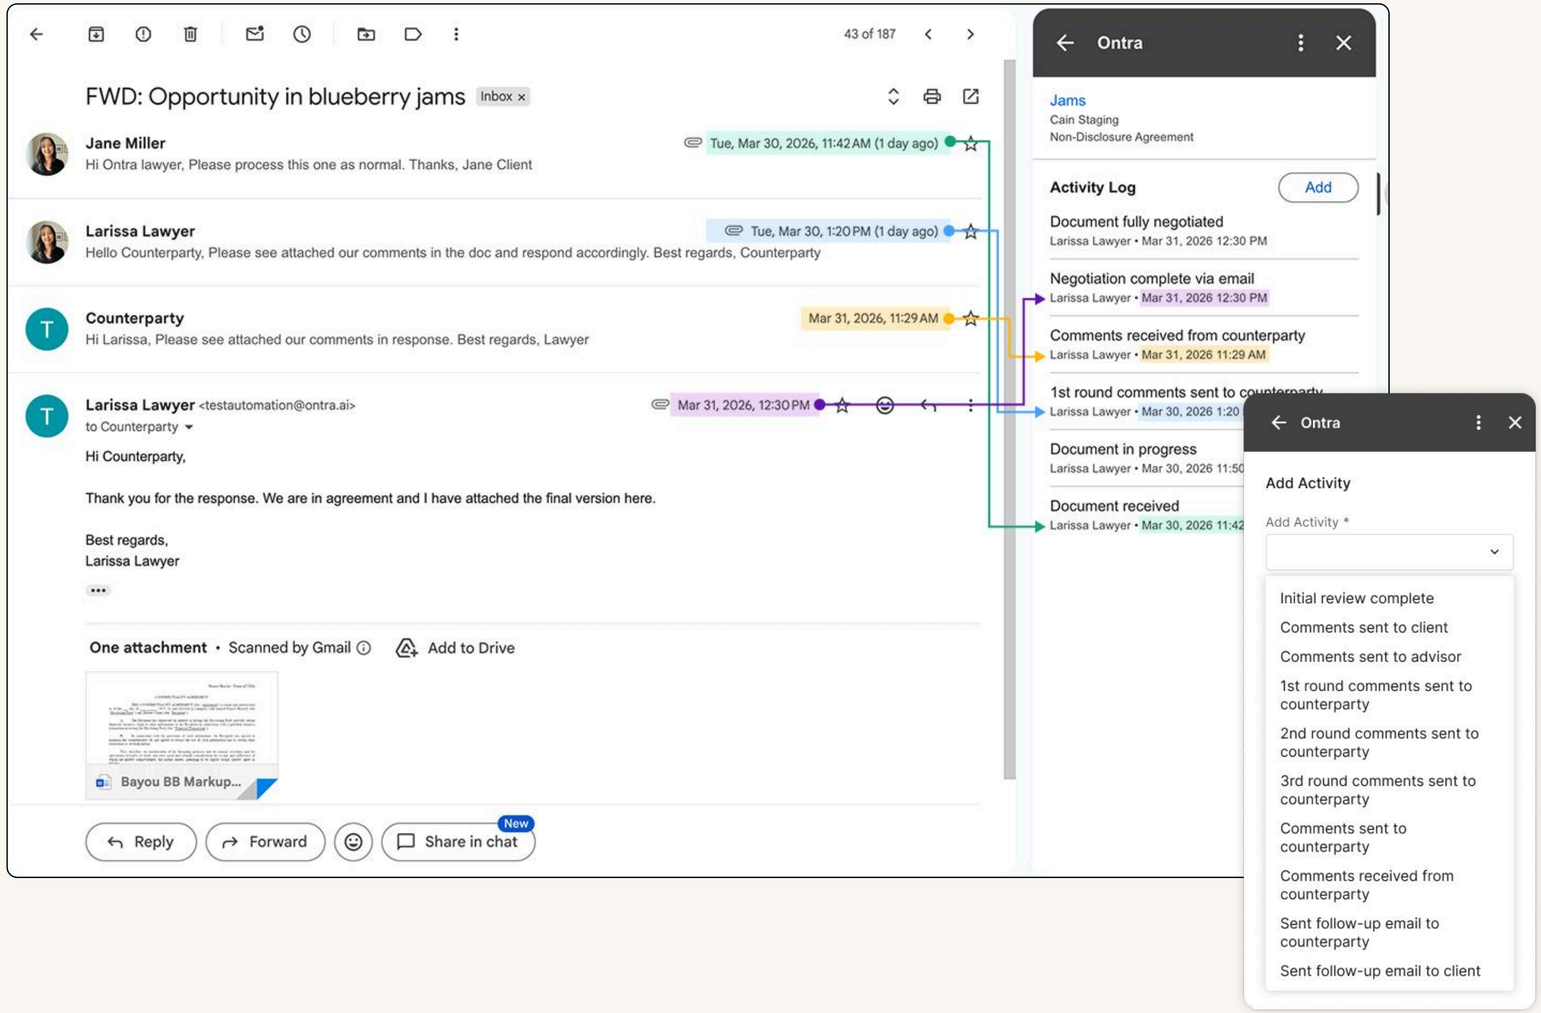
- Manually add activity updates that derive the timestamp from the email itself

CONSTRAINTS

- Existing list of activities in the web app
- Google’s UI components

ROLLOUT PLAN

- Pilot with a subset of the Ontra Legal Network (OLN) who are high performers, scaling up, and motivated to work with us
- Roll out to remaining OLN, then build in automation with AI



Challenges

RESULT

For the first time in Ontra history, we achieved fully verified document tracking.

But there was still a long way to go with adoption: our most enthusiastic users tracked ~25% of their documents, and the others tracked much fewer.

CONSTRAINTS

Latency was the #1 user complaint, but since we were limited by Google's building blocks, reducing latency was more difficult and we faced diminishing returns.

It was also tricky to deal with cases where customers sent multiple documents per email thread, rather than just one.

FEEDBACK

"It's nice to have the accurate timestamps, and I like that I don't have to do it right away...

But I don't know if it's saving me that much time."

DEVELOPMENT

Gmail Add-On v2: Classifying Messages

FUNCTIONALITY

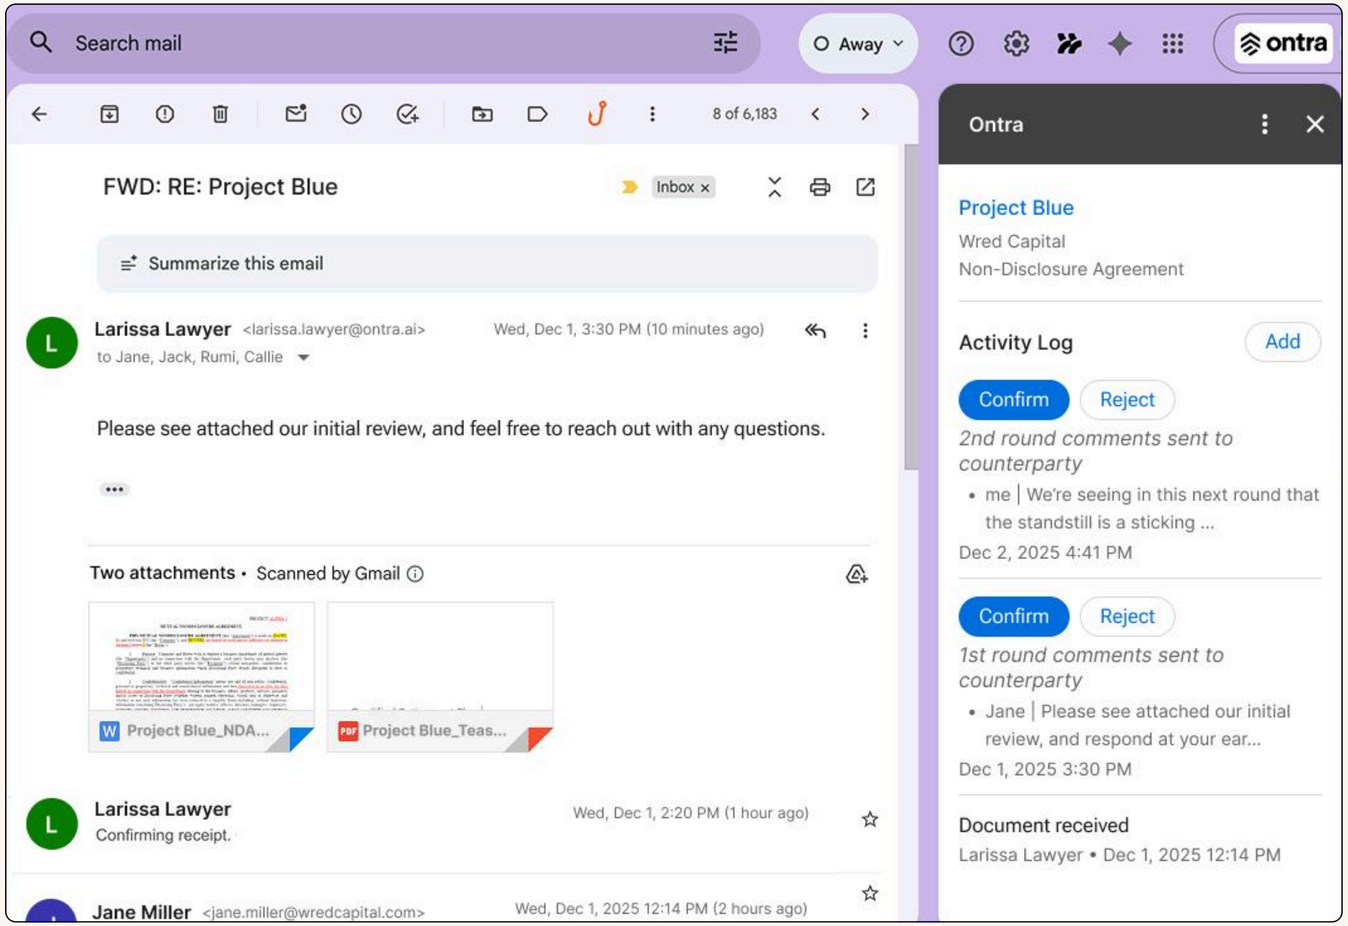
- Within verified email threads, AI suggests activities based on email content, which users can confirm or reject

CHALLENGES

- We already have access to shared inboxes, but customers often email individual lawyer emails instead
- Legal and security risks associated with accessing individual emails

FEEDBACK LOOP

- We worked even more closely with the internal Legal Network Success team to train users, gather feedback, and increase usage



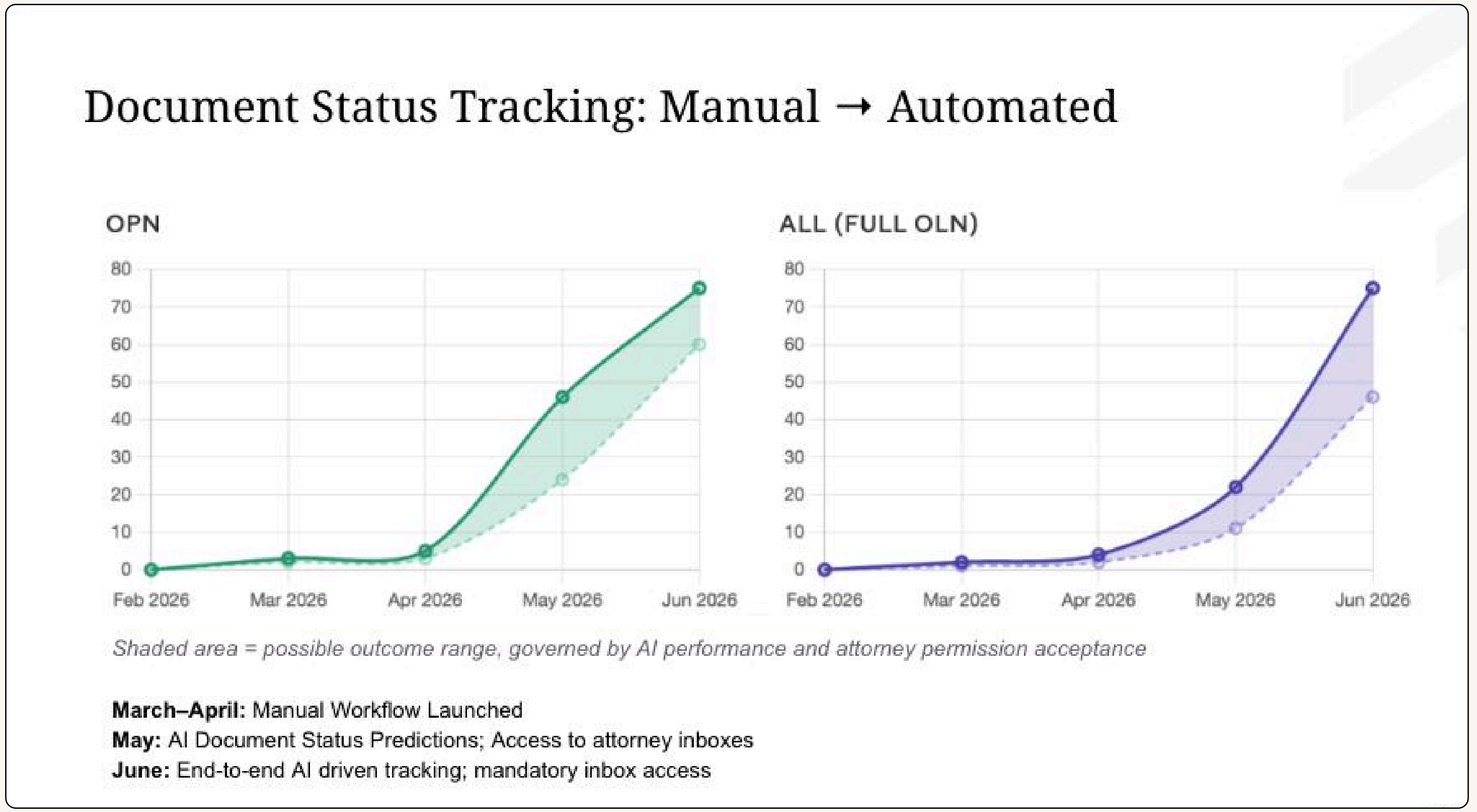
This bends the curve.

RESULTS

Within 2 days of launching this version to the entire OLN user base, we were able to start ingesting 5x more emails and generating predictions for them with a ~80% accuracy rate.

NEXT STEPS

Increase overall coverage by automatically identifying the correct email threads.



PROJECTIONS AS OF EARLY MAY

DEVELOPMENT (ONGOING)

Gmail Add-On v3: Classifying Threads

FUNCTIONALITY

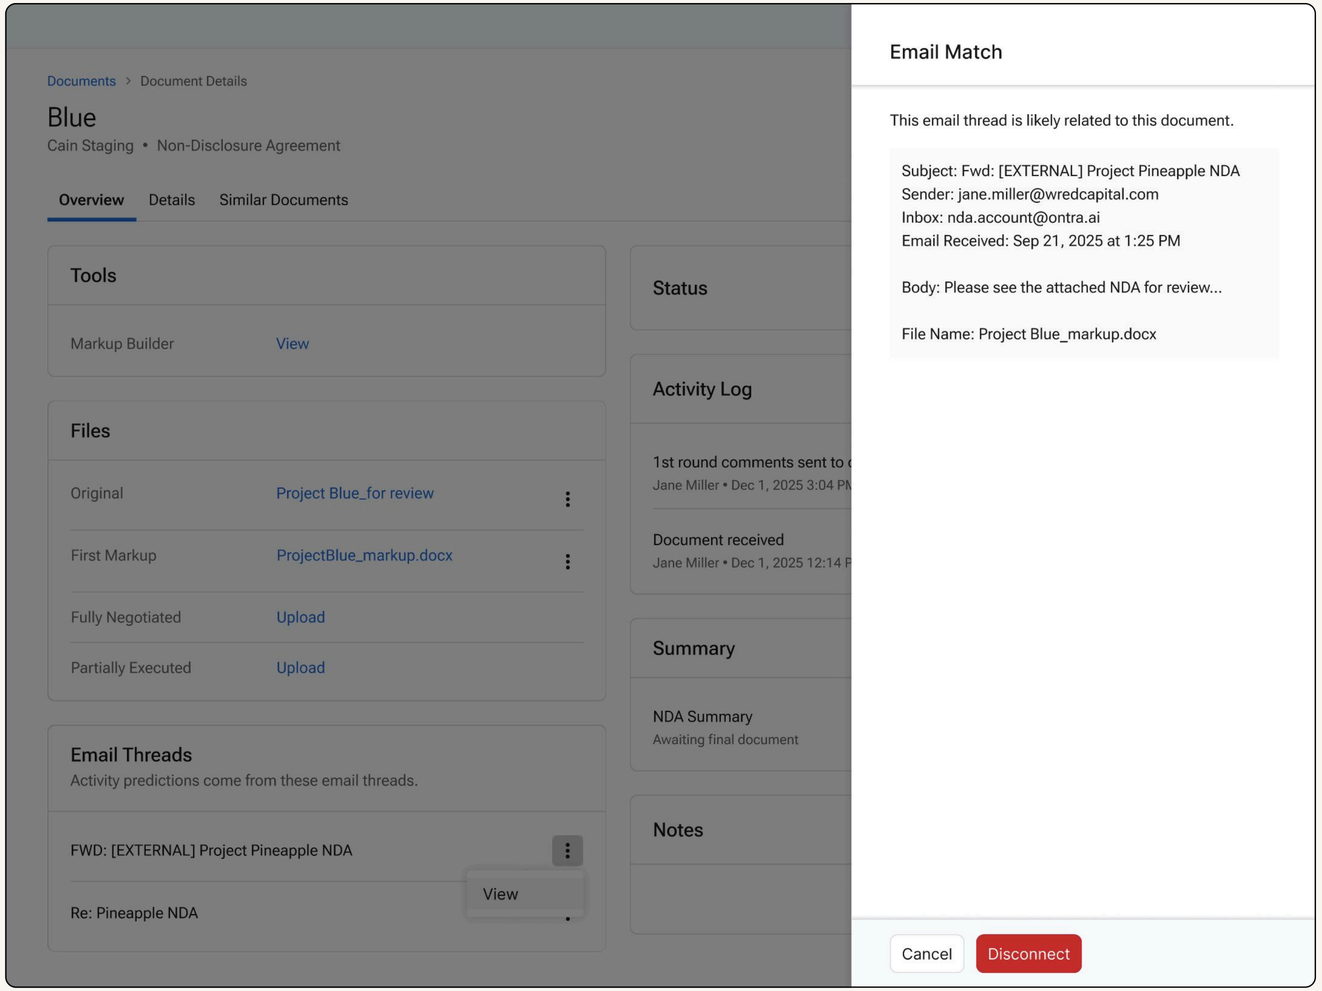
- We use a mix of heuristics and AI to identify whether an email thread is relevant to a document negotiation

CHALLENGES

- Reduced team capacity: I start covering 2 teams due to our PM being pulled for a tiger team; the other PM leaves the company
- AI engineering bottleneck

HYPOTHESIS

- This unlocks another chunk of document coverage, again without relying on human inputs



Results

SHIFTING THE CONVERSATION

By mid-June, we were well on track to gather the broadest dataset that Ontra had ever had before about the status of their documents.

We started to look farther ahead: what can other teams do with this data? What do we need to do to ensure that this data is used fruitfully?

Key Workflow Milestone (OPNs)	Priority	Start	May	June
Negotiation Started Client sends in doc	High	~48%	50.3%	71%
First Markup Sent Lawyer sends first redline to counterparty	High	0	34%	58%
Counterparty Markup Received CP responds with markup	Moderate	0	19.7%	27%
Subsequent Markup(s) Sent Lawyer sends additional markup(s) to counterparty	Moderate	0	14.8%	18.7%
Fully Negotiated Lawyer returns fully negotiated doc to customer	High	0	41%	32%
Executed Required party signs the doc	Moderate	0	38%	27%
Escalated Attorney contacts client for assistance	Moderate	0	19%	28%
Fully Tracked Has all high-priority milestones		0	20%	21%

ONLY INCLUDES JUNE 1-11, WILL INCREASE AS DOCUMENTS PROGRESS

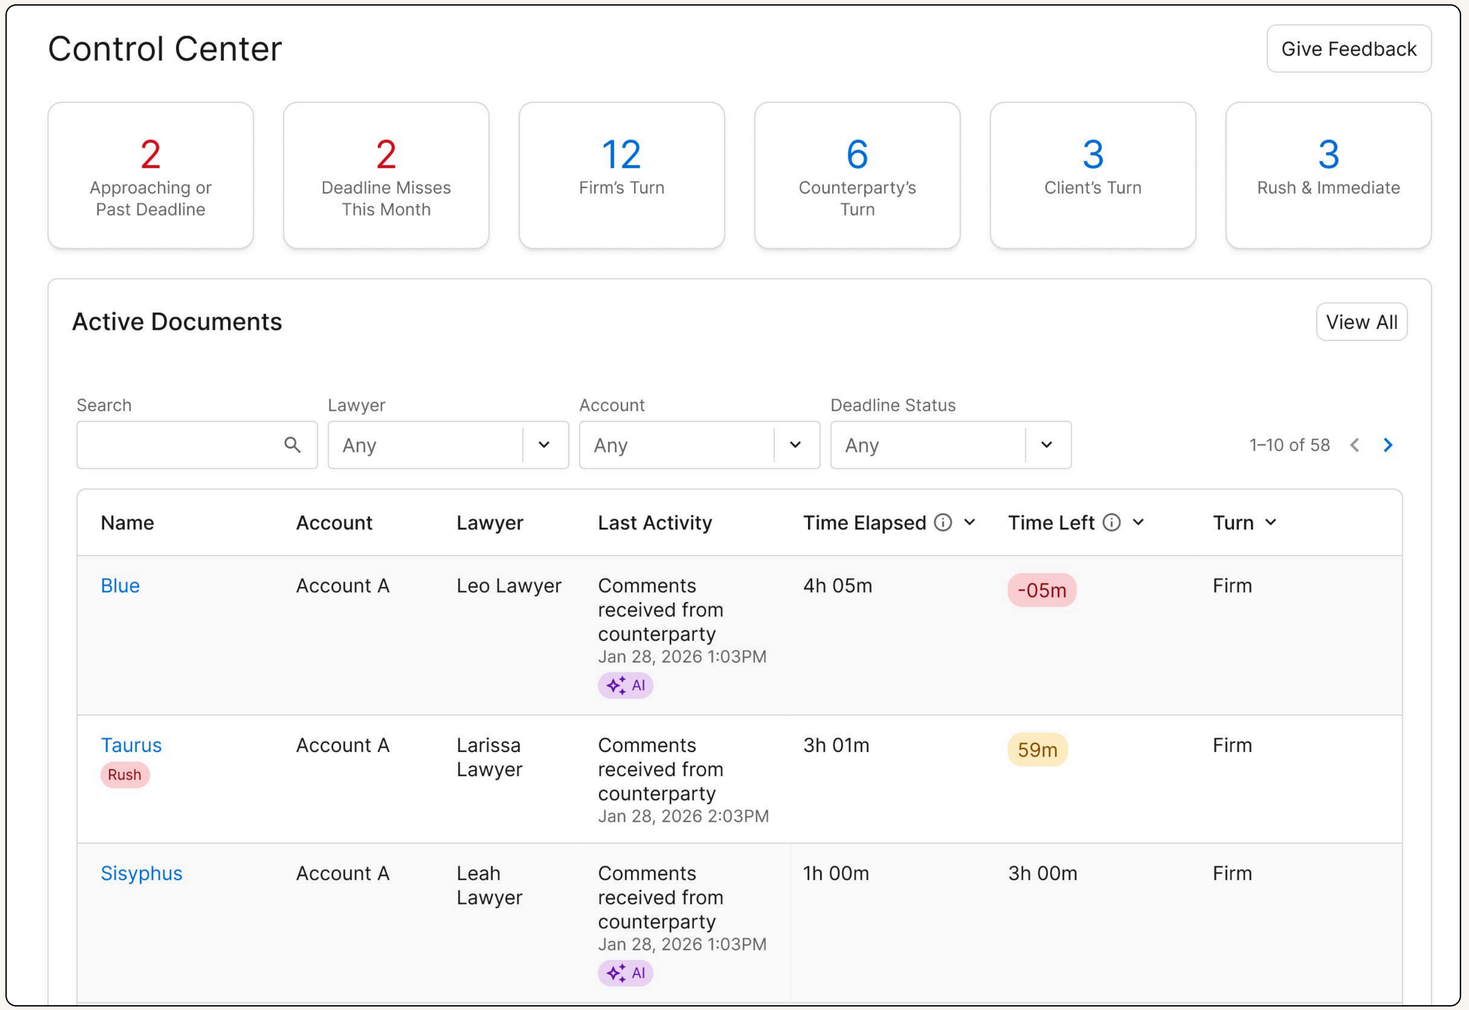
MODERATE PRIORITY ITEMS DO NOT OCCUR FOR EVERY DOCUMENT

Results

TARGETING BROADER USAGE

I worked with the sister team to incorporate the AI predictions more fully into the Ontra app, having them automatically calculate deadlines and track any missed ones.

Documents falling through the cracks was a perpetual fear of the Ontra lawyers, and this was the place to go to prevent that from happening.



Lessons Learned

USER BEHAVIOR

People can overestimate how likely they are to change their behaviors, especially with complex workflows that they've worked on refining for years.

ORG ALIGNMENT

AI-generated data is much more accepted by the Ontra Legal Network, but gaining alignment across the org internally about using that data for other purposes is its own challenge.

HUMANS IN THE LOOP

User feedback remained an important tool: by running feedback sessions, I learned where the AI predictions faltered and helped our team improve them.

What I Brought to the Team

UX Research

- Planned, scheduled, and facilitated all feedback sessions

UX Design

- Owned the designs for the Gmail Add-On across all phases

Product Strategy

- Contributed to defining success metrics for each quarter and milestone
- Worked with the PM and Engineering Manager to demo progress each week to C-suite

Cross-Team Enablement

- Created documentation, ran trainings, demo’ed features, and communicated new releases to other internal teams and to user base directly

Adjacent Team — Product & UX

- In late April, stepped in as designer for the adjacent team building the Control Center, where all document and timestamp data surfaces in one place for "air traffic control" visibility. Identified and advocated for fixes to better align data usage.