



LIVE WEBINAR · 16 SEPTEMBER 2026

The 2027 operating model for in-house legal

Building systems your team can actually rely on, and the
skills it takes to run them

Lorna Khemraz

LEAD COUNSEL & LEGAL AI ALIGNMENT, FLANK





Where we're going

- 1 Where we are
- 2 The operating model for in-house legal
- 3 How you build it
- 4 The system that helps it scale
- 5 In conversation with Alistair Maiden



LIVE POLL

How far would you trust an agent with one routine workflow, say NDAs, this quarter?

Answer in the poll panel. I'll ask the same question again at ten to five.

- 1 Not at all, not yet
- 2 Only if a lawyer checks every output
- 3 If exceptions are escalated and someone owns the playbook
- 4 Fully, for routine work within our rules
- 5 We already do this



The foundation is already there.

THE PROBLEM

Inexpensive work is being done by expensive resources.
resources.





Three things you already ask of a new hire

Capability

FOR A HIRE

Knowledge and competency. competency. Can they do the the job?

FOR AN AGENT

Skills and architecture. What is the agent built to do?

Context

FOR A HIRE

Organisational context so they they can operate in the role. role.

FOR AN AGENT

The playbook, written for a reader that can't ask what you you meant.

Scalability

FOR A HIRE

Grows into the role. Trusted to work semi-autonomously.

FOR AN AGENT

The system around the agent: agent: cockpit, escalations, data, an auditor.



Many teams buy the tech but skip the control room

- 1 Does the job well
The standard you'd expect of a person, every time.
- 2 Auditable
You can see what it did, and why.
- 3 Robust
Holds up on real paper, with real users.



HOW WE GOT HERE

Three eras of the legal stack

Pre-AI

TOOLS

Microsoft, a CLM – e.g iManage or iManage or HighQ - Carpe Diem Diem for your time.

Humans do the work.

AI-powered

FASTER TOOLS

Same stack plus Copilot, ChatGPT ChatGPT or Gemini. Same work, work, one prompt at a time, faster.

Humans do the work with AI.

Agentic

WORKFLOWS

Repeatable work runs end to end through agents that know your positions.

Agents do the work. Humans supervise.

64%

of in-house teams expect AI to reduce their reliance on outside counsel

7%

have seen total matter cost fall

ACC / Everlaw



WHAT ACTUALLY CHANGES

From prompting tasks to running workflows

AI-POWERED, TODAY

- Open Copilot. Paste the contract.
- Prompt. Read. Prompt again.
- Re-explain your context every time.
- Do the mark-up yourself.



AGENTIC, 2027

- Sales emails legal, as they always have.
- A triage agent routes it to the right agent. Nobody forwards anything.
- The work is done overnight, to your playbook.
- Waiting in your supervision queue by morning. Twenty minutes of judgment.



Is it a candidate for automation?

1

Is it novel every time?

No. A deal contract landing on your desk is rarely rarely new.

2

Is it repetitive?

Yes. The same clauses, the same objections, the objections, the same positions.

3

Would you give it to a junior for a first first pass?

Yes, if you had the team. That is exactly the shape of work an agent can take.

4

Can you articulate your approach?

Yes, or it can be inferred from your precedent red precedent red lines.




If you can explain it easily to a human, you can explain it to an agent.

RIPE FOR AN AGENT

- Drafting a full agreement pack from from a signed MOU, on your own own templates
- Handling counterparty objections on your own paper
- First-round review of third-party paper against your playbook
- Answering policy questions from the business
- Triaging the shared legal inbox



From red lines to a system you can rely on



Rules an agent can follow are rules a human can defend.



Where it can go wrong, and how to steer the agent

1

Misread the play

Applied it in a way its author never intended.

2

Clumsy execution

Knew what to do. Made an untidy untidy redline of it.

3


Escalated needlessly

Sent for further review what the the playbook had already answered.

These are the mistakes a bright junior makes in their second week. Same answer: supervision, feedback, a supervision, feedback, a tighter play. Nobody gives a second-week junior signature authority.



The whole arc, **not just the review step**





The cockpit: review by exception

1

Escalations first

High-risk terms, playbook-flagged items, anything needing a business owner's input. input.

2

Judgment on top

Override the playbook for this deal, this counterparty, this jurisdiction. Record why.

3

Human approves

The agent's review is for your information. Nothing leaves without a person clicking approve.

The screenshot shows the Flank* interface for reviewing a document titled "Lumen Logistics - Mutual NDA.docx". The main content area displays the document text, which includes definitions of terms like "Confidential Information", "Representatives", and "Affiliate", as well as sections for "Purpose", "Confidentiality Obligations", and "Confidentiality Term". The right sidebar shows a list of "Playbook checks 4" and "Flank flags 3". The checks include items like "Marking-independent definition", "Return/deletion timing window", "Professional adviser retention carve-out", "Confidentiality term", "Purpose tied to a defined opportunity", "Governing Law", "Governing law and jurisdiction", "Counterparty Policies", "Indemnities", "Equitable Remedies", and "Permissive 'may seek' framing". Each check has a status indicator (green for compliant, red for non-compliant). At the bottom, there is a progress bar and a button labeled "Apply all (4)".

Q ↶ ↷ 100% ▾ Heading 1 ▾ B I U ⋮ Reviewing ▾ ⬇

SHALL NOT EXCEED THE TOTAL SUBSCRIPTION FEES PAID BY CUSTOMER TO BLUESTONEHR DURING THE ~~TWELVE (12) MONTHS IMMEDIATELY PRECEDING TWENTY-FOUR (24) MONTHS IMMEDIATELY PRECEDING~~ THE DATE OF THE CLAIM. ~~CUSTOMER'S LIABILITY TO BLUESTONEHR SHALL BE UNCAPPED AND SHALL INCLUDE ALL DIRECT, INDIRECT, SPECIAL, AND CONSEQUENTIAL DAMAGES.~~ THE FOREGOING CAP SHALL APPLY EQUALLY TO THE LIABILITY OF EACH PARTY. THE CAP SHALL NOT APPLY TO LIABILITY ARISING FROM A PARTY'S BREACH OF ITS CONFIDENTIALITY OR DATA PROTECTION OBLIGATIONS, OR ITS GROSS NEGLIGENCE OR WILFUL MISCONDUCT.

- 9.3** Exceptions. The liability cap in Section 9.2 does not apply to Customer's payment obligations or Customer's indemnification obligations under Section 8.1. Nothing in this Agreement limits either party's liability for death or personal injury caused by negligence, or fraud.
- 9.4** Essential Basis. The parties acknowledge that the limitations of liability in this Section 9 reflect a reasonable allocation of risk and are an essential element of the basis of the bargain between the parties.

10. TERM AND TERMINATION

- 10.1** Term. This Agreement commences on the Effective Date and continues for the Initial Subscription Period set out in Schedule 1 (the "Initial Term"). Thereafter, this Agreement shall automatically renew for successive periods of 12 months (each a "Renewal Period", and together with the Initial Term, the "Term") unless either party provides written notice of non-renewal ~~at least 90 days prior to the end of the then-current period~~ at least 30 days prior to the end of the then-current period.
- 10.2** Termination for Cause. Either party may terminate this Agreement immediately upon written notice if the other party: (a) materially breaches this Agreement and fails to cure such breach within 30 days of receiving written notice specifying the breach; or (b) becomes insolvent, makes an assignment for the benefit of creditors, or is the subject of a bankruptcy, administration, or liquidation proceeding.
- 10.3** Termination for Convenience. **BlueStone** may terminate this Agreement for any reason upon 30 days' written notice to Customer. **Customer may not terminate this Agreement for convenience during the Initial Term.** Following the Initial Term, Customer may terminate for convenience upon ~~120 days' prior written notice~~ 60 days' prior written notice.

Third-party IP indemnity

Third-party IP indemnity

Non-compliant

Customer data & AI training

Customer data & AI training

Non-compliant

Service levels & credits

Service levels & credits

Non-compliant

Term & termination

Term & auto-renewal

Non-compliant

Shortened non-renewal window to 30 days and convenience notice to 60

Playbook standard >

2 applied changes

Non-renewal notice — 30 days

Revert

Change applied

This Agreement shall automatically renew for successive periods of 12 months unless either party provides written notice of non-renewal at least 99 30 days prior to the end of the then-current period.

Comment — sent with the redline

A 90-day non-renewal notice makes the exit easy to miss. Cut to 30 days.

\$10.1 ↗

> Termination for convenience — 60 days

Revert (2)

Approve changes (2)

Inbound

117

15 filtered · 101 real requests

Agent-handled

49%

46 auto · 3 supervised

Median agent resolution

6m

90% resolve within 1.6h

Cost displaced


£26.8k

≈ 107h · £108.8k annualised


Of 117 inbound, 15 filtered as noise. 49 (49%) handled end-to-end by agents — median resolution 6m, 90% inside 1.6h. 52 routed to humans, who take 20h (median) just to first respond — and there visibility ends. £26.8k of internal-lawyer time displaced this window (≈ 107h at the blended rate).

Request flow

Where 117 inbound requests went, 90 days.



Weekly volume



Agent Human Unrouted

Demand by use case

Agent auto Supervised Internal External Delivery centre Unrouted

Click a row to see who handles it, where the bottlenecks are, and which clauses are driving cycle time.

Use case	Handled by	Median resolution	90% within	Status
> Policy 26 matters		5m	6m	Bottleneck
> Franchise agreement 21 matters		29m	1.6h	—
> MSA review 12 matters		2.6h	2.6h	Bottleneck

⚙️ Across **2 active playbooks**, 8 clauses are being observed in counterparty traffic. **2 flagged to tighten** (position or fallback isn't landing), **2 candidates to relax** (fallback lands almost every time — promote it), 4 holding steady. Open a playbook to see clause-level pushback and acceptance data.

▼ NDA

Northwind paper · drafting + counterparty review

5 clauses 1 tighten 1 relax 3 hold

Clause	Pushback	Accepted as is	Fallback accepted	Redlines	Flank flag
> Definition of confidential information	88%	12%	38%	14	Tighten
> Term length & survival period	71%	29%	92%	12	Relax
> Indemnity scope	64%	36%	74%	11	Hold
> Governing law	54%	46%	81%	9	Hold
> Non-solicit	22%	78%	96%	8	Hold

Based on counterparty redlines observed in the last 90 days.

Open in playbook editor ➤

▼ Franchise agreement

Northwind paper · drafting + counterparty review

3 clauses 1 tighten 1 relax 1 hold

Clause	Pushback	Accepted as is	Fallback accepted	Redlines	Flank flag
> Territory exclusivity expansion	94%	6%	23%	10	Tighten
> Marketing fund contribution	33%	67%	95%	7	Relax
> Royalty step-down	47%	53%	88%	10	Hold

- + New session
- 📧 Inbox
- 🔍 Intelligence
- 🛠 Customize

🔍 **8 items need attention today** — 1 unrouted, 3 waiting on supervisor review, 4 aged past 24 hours. 2 use cases concentrated on one person. Agents handled 49 matters end-to-end this window; 52 were routed to humans.

Needs attention

Click through to act.

🕒 Unrouted

1

No routing rule matched. Open in inbox to triage.

👤 Awaiting supervisor

3

Agent-drafted, ready for a lawyer to review and send.

🕒 Aged past 24h

4

Open requests past a working day. Assignees flagged below.

Request	Assignee	Category	Age
Property notice — Stonecrust multi-site closure	Marcus Webb	Fallback	38h
MSA — BridgePoint Wholesale, third revision	Marcus Webb	MSA review	32h
Sanctions screening — new APAC franchisee	Sarah Chen	Regulatory	27h
Influencer partnership terms — Heath & Spoon	—	Fallback	24h

What's causing delays

Structural causes — not symptoms. Fix these to clear the queue at the source.

- Requests with no routing rule

1 this window

Three recent requests fell to the GC inbox because no rule matched (brand collaborations, property leases, influencer partnerships). Each sits ~24h before a human triages it.

Close gap from inbox
- Marcus is at capacity on MSAs + procurement

100% utilisation

Marcus is handling 20 routed items (MSA review + procurement). Average pickup is 19.3h; 16 cleared, oldest open is 862h.

Review automation candidate
- Regulatory and HR escalations all land on the GC

Avg pickup 19.5h

🚀 **2 work types** are ready to scope today, with 1 more close behind. Each card below shows the work type, automation fit, recommended agent, and the caveats to think through before kicking off.

● **Ready** 2 candidates

Policy

Ready

Internal policy questions on contracting limits, approvals, supplier confidentiality.

Volume / mo	Annual savings	Repetitiveness
19	£18.6k	92%

Automation fit 75

Recommended

Policy Q&A agent

Caveats

- Confidence threshold tuned high to prevent hallucinated policy positions
- Falls back to a human when policy hasn't been updated in >18 months

Procurement

Ready

Vendor onboarding paperwork — order forms, DPAs, SOWs against standard positions.

Volume / mo	Annual savings	Repetitiveness
7	£6.8k	86%

Automation fit 51

Recommended

Vendor & procurement agent

Caveats

- Limited to standard order forms and DPAs — non-standard flows still routed to Marcus
- Requires connector to procurement system for matter-of-record creation

● **Almost ready** 1 candidate

MSA review

Almost ready

Counterparty MSAs reviewed against the master playbook — liability cap, indemnity, IP, termination.

Volume / mo	Annual savings	Repetitiveness
9	£8.4k	78%

Automation fit 51

Recommended



From firefighting to a legal engine

Someone is the auditor in charge. For the first time, they can actually see the function.

Edit patterns

What humans keep changing tells you what what the playbook should say.

Pushback patterns

Which positions counterparties resist, and how often you concede.

Query flow

What lands, from where, where, and who picks it it up.

Resourcing

Budget and headcount decisions made on evidence, not on who shouted loudest.



SOME COMMON CONCERNS

Three questions I get every time

"Who's responsible if it gets it gets it wrong?"

The human in the loop is there by there by design. Playbooks are are signed off by legal, escalations escalations route to legal, and nothing reaches a counterparty or counterparty or signature without without a person clicking approve. approve.

"What does this mean for my my team?"

An expansion of capacity rather rather than a like-for-like swap. swap. Your people run agents that that are extensions of them, and and their job moves up a level. level.

"Aren't we creating a dependency?"

The playbook is your judgment, judgment, written down and owned by you. You choose who who supervises: your own lawyers, a partner firm, or Flank's Flank's team. The expertise stays stays in the building – you're insourcing.



What the senior associate of agents **actually** does

1

Articulate judgment

Turn 'how we do things' into a playbook: positions, exceptions, escalation triggers.

2

Review by exception

Work from escalations, not every line. Override
Override with a reason, and record it.

3

Read the data

Treat edit patterns, pushback and query volumes
volumes as management information.

4

Own the audit

Someone is accountable for agent performance,
performance, the way someone supervises a
a trainee.



Keep a supervision matrix

1 Name an owner and a backup

These individual hold admin rights.

2 Agree who owns playbook hygiene

Monthly to start with, then quarterly.

3 Review the access list regularly.

Weekly in pilot. Re-check monthly + on every joiner and leaver.

Person	Platform role	NDA agent (US)	NDA agent (CA)	NDA agent (MX)
A. Rivera (legal ops)	App admin	Admin	Admin	Admin
J. Doe (US)	Agent admin	Admin	User	User
J. Smith (US)	Agent admin	Admin	None	None
M. Tremblay (CA)	Agent admin	None	Admin	None
C. Hernández (MX)	Agent admin	None	None	Admin
P. Patel (IT)	Chat user	User	User	User
S. Lee (procurement)	Chat user	None	None	User
Owner / backup		Doe / Smith	Tremblay / Doe	Hernández / Doe
Company access		On	On	Off
Last reviewed		1 Sep 2026	1 Sep 2026	1 Sep 2026

Dummy data.



Rebuild the playbook from 10 to 20 red lines

1

Pull the mark-ups

One contract type. The versions matching the level of level of negotiation you want to automate, e.g. first first round.

2

List every position you took

Compare with the playbook. Every gap is an unwritten unwritten decision.

3

Write each gap as a play

Position, fallback, threshold, escalation.

Position	Playbook says	Red lines show	The play
Non-solicit period	6 months	12 months	12. Escalate above 18.
Liability cap	Silent	2x fees under 50k	2x under 50k. Escalate above.
Governing law	England	Ireland, 3 times	England. Ireland with sign-off.
Assignment	Consent required	Never withheld	Not unreasonably withheld.
Payment terms	30 days	45 for two accounts	30. 45 for listed accounts.

Illustrative.



Run the four questions on one workflow

- 1** Pick the work that lands most often
Not the most interesting. The most repeated.
- 2** Ask the four questions honestly
They're on the right.
- 3** No, yes, yes, yes
That is your first candidate. One workflow, supervised, supervised, a few weeks.

WORKED EXAMPLE: A DEAL CONTRACT

Is it novel every time? No

Is it repetitive? Yes

Would a junior take the first pass?
pass? Yes

Can you say how you approach it? Yes

If you can explain it easily to a human, you can explain it to an agent.



LIVE POLL, SAME QUESTION

How far would you trust an agent with one routine workflow, say NDAs, this quarter?

Forty minutes on. Has anything moved?

- 1 Not at all, not yet
- 2 Only if a lawyer checks every output
- 3 If exceptions are escalated and someone owns the playbook
- 4 Fully, for routine work within our rules
- 5 We already do this



You already know how to do this.

Capability

Agents built for the work

Context

Playbooks that hold your judgment

Scalability

A cockpit, escalations, data, and a lawyer as auditor

Agents in the seats, lawyers in the cockpit, and the same three questions you have always always asked of a new hire.



What's holding legal leaders back, and what tips them over?

Alistair Maiden

FLANK

Flank*





Any questions?

Add us on LinkedIn.

We'd love your thoughts on this session so we can shape the next one.