

HOW TO TEST AI WITHOUT BREAKING THE BILLABLE HOUR

The Hybrid AI Firm Playbook

A partner-safe framework for borrowing AI-native efficiency without blowing up pricing, training, governance, or trust.

INSIDE

- Hybrid Firm Map
- Hidden billable-hour economics
- Workflow readiness scorecard
- Hybrid pricing ladder
- 30-day pilot plan
- Partner meeting checklist

Manu Ayala · AI Vortex

For small, mid-size, and boutique law firms testing AI without pretending the billable hour disappears overnight.

READ THIS FIRST

EXECUTIVE SUMMARY

This playbook is for small, mid-size, and boutique law firms that know AI matters but do not want to pretend the billable hour disappears overnight.

Most AI advice skips the business-model problem. It tells firms to buy a tool, train the team, or copy AI-native firms. That misses the harder reality: traditional firms already have pricing models, training models, client expectations, partner economics, confidentiality constraints, and workflows that evolved over years.

The practical path is not to become AI-native overnight. The practical path is hybrid: borrow the operating logic that compounds, test it in one workflow, measure the result, and decide what can safely change.

**The question is not whether AI can save time.
The question is where that saved time goes.**

CATEGORY

WHAT IS A HYBRID AI FIRM?

Traditional firms are not becoming AI-native overnight. But they are also not staying purely traditional. The practical path is hybrid.

The goal is not to turn every matter into a product. The goal is to identify which parts of the firm can safely become more repeatable, measurable, governed, and efficient without weakening the parts clients actually trust the firm for.

THE REAL PROBLEM

A Hybrid AI Firm borrows AI-native operating logic without pretending every workflow should become automated, fixed-fee, or AI-native overnight. It is workflow-first and partner-safe.

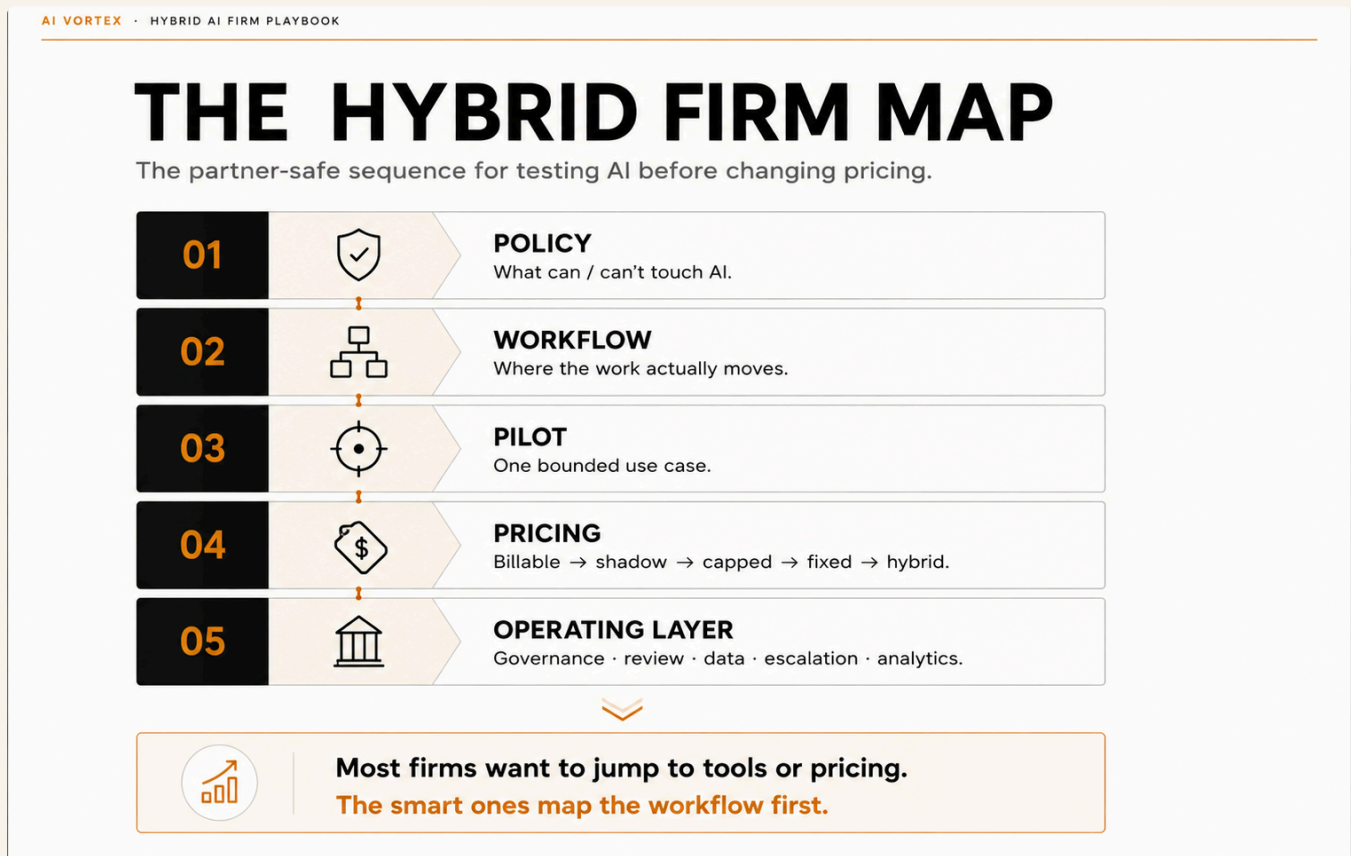
- **Not anti-billable:** billable work can remain billable.
- **Not AI-native theater:** no performative "AI transformation."
- **Not vendor-first:** the tool is downstream of the workflow.
- **Workflow-first:** map the work and the drag before automation.
- **Partner-safe:** evidence-based tests before pricing changes.

The safest path is not to jump straight to a tool, a pricing model, or a full operating layer. The sequence matters.

FRAMEWORK

THE HYBRID FIRM MAP

Use this as the operating sequence. Policy before workflow. Workflow before pilot. Pilot before pricing. Pricing before scale.



HOW TO USE THIS MAP

- Do not treat it as a maturity badge. Treat it as a sequence.
- Policy answers what can safely touch AI.
- Workflow shows where the work actually moves.
- Pilot creates evidence before pricing changes.
- Operating layer comes after the firm knows what to scale.

ECONOMICS

THE HIDDEN ECONOMICS OF THE BILLABLE HOUR

The billable hour is not only a pricing model. It is a measurement system, training structure, incentive structure, utilization model, and revenue architecture. That is why “AI saves time” becomes complicated inside a traditional firm.

A matter may generate ten billable hours, but the real economics also include the work around the work: chasing documents, cleaning intake, updating status, reconstructing context, creating internal notes, passing information between systems, managing versions, and repairing templates.

THE INVISIBLE WORK TAX

The Invisible Work Tax is the non-billable drag around legal work that quietly reduces margin, capacity, focus, speed and quality of life without appearing as a clean line item.

- Searching for documents and prior versions
- Cleaning messy client intake
- Drafting repeated follow-up emails
- Updating matter status manually
- Rebuilding context after handoffs
- Maintaining spreadsheets, trackers and templates
- Writing internal summaries nobody reuses

VALUE CAPTURE

NOMINAL HOUR VS REAL HOUR




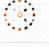








The billable hour is visible. The friction around it is often invisible.


AI VORTEX · HYBRID AI FIRM PLAYBOOK


ECONOMICS

THE NOMINAL HOUR VS. THE REAL HOUR

Your hourly rate is nominal. Your workflow determines the real economics.

 NOMINAL HOUR	VS.	 REAL HOUR
 Billing rate		 Invisible work tax
 Billable time		 Intake drag
 Visible revenue		 Document chasing
 Utilization		 Status work
		 Handoffs
		 Context rebuilding

 **Real economics** = **Billable work** - **Invisible work tax**

 The first AI win doesn't need to reduce billable hours. It can reduce the invisible work tax around them.

This is why the first AI win does not have to reduce billable hours. It can recover capacity around the billable work.

**The first AI win does not need to reduce billable hours.
It can reduce the invisible work tax around them.**

PRESSURE SIGNAL

AI VALUE IS NOT ONE THING

A traditional firm does not need to capture AI value in one way. The right pilot depends on what the firm wants to recover.

Saved time can become margin, capacity, risk control, training leverage, client clarity, or recovered time.


AI VORTEX · HYBRID AI FIRM PLAYBOOK

VALUE CAPTURE

AI VALUE IS NOT ONE THING

Saved time can become different kinds of value depending on the firm.


01



CAPACITY

Handle more work or higher-value matters with the same team.


02



MARGIN

Reduce non-billable drag around matters and protect realized margin.


03



RISK CONTROL

Create more consistent review and escalation before issues escalate.


04



CLIENT CLARITY

Improve expectations, status updates, and turnaround communication.


05



TRAINING LEVERAGE


Turn review into reusable learning for lawyers, juniors, and the firm.

06



RECOVERED TIME

Reduce nights and weekends spent on low-value, manual work.



Before changing pricing, decide what saved time is for.

The key decision is not simply whether AI saves time. The key decision is whether saved time becomes capacity, margin, risk control, training leverage, client clarity, or recovered time.

ECONOMICS

WHAT AI-NATIVE FIRMS GET RIGHT

AI-native firms are not interesting because every traditional firm should copy their surface. They are interesting because they show what happens when workflow, pricing, review, and delivery are designed together from the beginning.

Many AI-native models start with bounded work, clearer scopes, flat or packaged pricing, faster turnaround, and lawyers focused on review and judgment. That operating logic matters. But copying the price or the marketing is not the same as understanding the system underneath.

FOLLOW THE MONEY BEFORE COPYING THE MODEL

A traditional firm should not copy the price of a VC-backed AI-native firm unless it understands the economics subsidizing that price. Some AI-native firms may be optimizing for growth, narrative, data, market share, volume, product refinement or future margin assumptions. A traditional firm may be optimizing for profitability today, partner distributions, quality, reputation, client relationships and sustainable risk.

The retail analogy is useful. A physical store should not copy every e-commerce price just because e-commerce has lower floor costs. It should study what digital operating logic can improve its own model without destroying what makes the traditional model valuable.

The point is not to copy AI-native firms. The point is to understand which parts of their operating logic can be safely adapted.

PRESSURE SIGNAL

BORROW THE OPERATING LOGIC

Traditional firms should study the logic underneath, not copy the surface.


AI VORTEX · HYBRID AI FIRM PLAYBOOK

PRESSURE SIGNAL

BORROW THE OPERATING LOGIC, NOT THE SURFACE

Traditional firms should study the logic underneath, not copy the surface.

❌ DON'T JUST COPY	🔄 FOCUS ON LOGIC	✅ STUDY THE LOGIC
<ul style="list-style-type: none"> Flat fee headline AI-native branding Low price Speed promise Vendor stack 		<ul style="list-style-type: none"> Scope control Review layer Exception rules Template discipline Workflow design Pricing confidence


AI-native firms are not interesting because of their surface.
They are interesting because of their system.

Traditional firms should not copy the surface: the price, the label, the speed promise, or the vendor stack. They should study the logic: scope control, review layers, exception rules, templates, workflow discipline, and pricing confidence.

CONSTRAINTS

WHY TRADITIONAL FIRMS CANNOT JUST COPY THEM

Traditional firms inherit economics and expectations:

- Billable hour economics and partner compensation
- Training structure and review loops
- Client trust and confidentiality obligations
- Inherited workflows and legacy toolchains
- Existing client expectations and delivery norms

THE BILLABLE HOUR IS NOT ONLY PRICING

- A revenue architecture
- A utilization model
- A partner compensation signal
- A training structure
- A client expectation model

The mistake is assuming a traditional firm can copy AI-native pricing before it has AI-native infrastructure.

That is why the transition has to be hybrid. The firm needs a way to capture AI efficiency without forcing the entire institution to change at once.

WHERE THE WORK MOVES

LEVEL 1: POLICY

Policy is not the whole strategy, but it is the first guardrail. A firm should not test AI workflows without deciding what can and cannot touch AI.

Policy should not be written as a generic AI disclaimer. It should answer operational questions: what data can be used, which tools are approved, who reviews outputs, when escalation is required, and what cannot enter an AI system at all.

MINIMUM POLICY DECISIONS

- Approved tools
- Banned uses
- Client data boundaries
- Anonymization rules
- Review standards
- Escalation owner
- Storage and audit trail
- Client consent triggers

**A public chatbot subscription is not an AI strategy.
Paid does not automatically mean governed.**

LEVEL 2

LEVEL 2: WORKFLOW

Most firms are not ready for a tool decision until they have made a workflow decision. The firm should map how work actually moves before deciding what AI should do.

Map one workflow from beginning to end: intake, document collection, cleanup, drafting, review, delivery, follow-up, billing, and post-matter learning. Then ask where drag appears and where judgment actually happens.

- Where does the work enter?
- Where does information get cleaned?
- Where is context reconstructed?
- Where do templates get reused or repaired?
- Where does a lawyer review?
- Where does judgment happen?
- Where does risk require escalation?
- Where does the client need clarity?

TOOL WARNING

A WRAPPER IS NOT AN OPERATING LAYER

A wrapper can be useful. A legal AI platform can be useful. A custom tool can be useful. The problem is buying one before deciding what workflow, review standard, data boundary, and escalation path it is supposed to support.

**A wrapper is not an operating layer.
If the workflow is unmapped, the tool will inherit the mess.**

Most firms do not have an AI problem yet. They have an unmapped workflow problem.

FRAMEWORK

LEVEL 3: PILOT

The first pilot should not be the flashiest use case. It should be the safest repeatable workflow with enough structure to learn from.

APPLIED MODEL

WORKFLOW READINESS SCORECARD

Before choosing a tool or changing pricing, score the workflow itself.

Score the workflow before you automate, price, or scale it.

FACTOR	QUESTION	SCORE (1-5) 1 = Low 5 = High				
01 REPEATABILITY	Does the work follow a repeatable pattern?	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
02 INPUT PREDICTABILITY	Do similar matters start with similar inputs?	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
03 SCOPE CLARITY	Can the output be defined before work starts?	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
04 SENSITIVITY / RISK	Can the work be tested safely?	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
05 REVIEWABILITY	Can a lawyer verify the output efficiently?	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
06 EXCEPTION VISIBILITY	Do you know what breaks the workflow?	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
07 CLIENT VALUE	Would speed or clarity matter to the client?	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
08 INTERNAL TIME LEAKAGE	Is there drag around the billable work?	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
09 GOVERNANCE READINESS	Are tool, data, and review rules clear?	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
10 PRICING POTENTIAL	Could this become capped or fixed later?	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
11 TRAINING IMPACT	Can learning improve, not disappear?	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

SCORE BANDS

45-55
Strong pilot candidate

35-44
Possible pilot, needs scoping

25-34
Workflow cleanup first

< 25
Not ready for AI-assisted pilot

The first benchmark that matters is your own baseline.

A low score does not mean AI is useless. It means the workflow needs cleanup before automation, pricing changes, or scale.

JUDGMENT LAYER

EXAMPLE WORKFLOW: RECORD SEALING / EXPUNGEMENT

This is a fictionalized example based on a common pattern: a firm has a traditional practice, but one workflow inside it is bounded enough to package, improve, and potentially price differently.

Use examples like this to separate the workflow layer (repeatable steps) from the judgment layer (lawyer review, risk calls, and client-specific nuance).

CURRENT WORKFLOW

- Client intake
- Eligibility facts
- Document collection
- Checklist or draft preparation
- Lawyer review
- Filing / delivery
- Status follow-up

POSSIBLE AI ROLE

- Organize intake
- Flag missing information
- Summarize client facts
- Generate an internal checklist
- Draft a status update
- Compare inputs against standard eligibility criteria

HUMAN ROLE

- Eligibility judgment
- Legal review
- Filing strategy
- Client advice
- Final approval

ESCALATION TRIGGERS

- Missing documents
- Prior denial
- Jurisdiction mismatch
- Unusual facts
- Urgent deadline
- Criminal history complexity

Pricing implication: this workflow may be a candidate for capped, fixed, or hybrid pricing only after exception rate and lawyer review time are measured.

PRICING BOUNDARIES

WHAT STAYS BILLABLE VS. WHAT CAN BE TESTED

Do not change pricing by practice area. Start by looking at the work pattern. The billable hour still makes sense where uncertainty is real and the judgment layer is doing the work.

Use the framework below to protect premium judgment work while you test repeatable workflow work in a bounded way.

AI VORTEX · HYBRID AI FIRM PLAYBOOK

PRICING / WORK PATTERNS

WHAT STAYS BILLABLE VS. WHAT CAN BE TESTED

Don't price by practice area. Price by work pattern.

<div style="background-color: #1a202c; color: white; padding: 5px; border-radius: 5px; margin-bottom: 10px;"> <div style="display: flex; align-items: center;"> <div> <p style="margin: 0;">STAYS BILLABLE / PREMIUM</p> <p style="margin: 0; font-size: 0.8em;">Uncertain, high-risk, or bespoke work.</p> </div> </div> <ul style="list-style-type: none"> <li style="border-bottom: 1px solid #ccc; padding: 5px 0 5px 20px; display: flex; align-items: center;"> Messy facts <li style="border-bottom: 1px solid #ccc; padding: 5px 0 5px 20px; display: flex; align-items: center;"> Uncertain scope <li style="border-bottom: 1px solid #ccc; padding: 5px 0 5px 20px; display: flex; align-items: center;"> Negotiation <li style="border-bottom: 1px solid #ccc; padding: 5px 0 5px 20px; display: flex; align-items: center;"> Litigation strategy <li style="border-bottom: 1px solid #ccc; padding: 5px 0 5px 20px; display: flex; align-items: center;"> High-risk judgment <li style="border-bottom: 1px solid #ccc; padding: 5px 0 5px 20px; display: flex; align-items: center;"> Bespoke advice <li style="padding: 5px 0 5px 20px; display: flex; align-items: center;"> Counterparty unpredictability </div>	<div style="background-color: #f59e00; color: white; padding: 5px; border-radius: 5px; margin-bottom: 10px;"> <div style="display: flex; align-items: center;"> <div> <p style="margin: 0;">CAN BE TESTED DIFFERENTLY</p> <p style="margin: 0; font-size: 0.8em;">Repeatable, structured, lower-risk work.</p> </div> </div> <ul style="list-style-type: none"> <li style="border-bottom: 1px solid #ccc; padding: 5px 0 5px 20px; display: flex; align-items: center;"> Repeatable intake <li style="border-bottom: 1px solid #ccc; padding: 5px 0 5px 20px; display: flex; align-items: center;"> Standard summaries <li style="border-bottom: 1px solid #ccc; padding: 5px 0 5px 20px; display: flex; align-items: center;"> First-pass document review <li style="border-bottom: 1px solid #ccc; padding: 5px 0 5px 20px; display: flex; align-items: center;"> Bounded contract packages <li style="border-bottom: 1px solid #ccc; padding: 5px 0 5px 20px; display: flex; align-items: center;"> Routine checklists <li style="border-bottom: 1px solid #ccc; padding: 5px 0 5px 20px; display: flex; align-items: center;"> Structured drafting <li style="border-bottom: 1px solid #ccc; padding: 5px 0 5px 20px; display: flex; align-items: center;"> Document comparison <li style="padding: 5px 0 5px 20px; display: flex; align-items: center;"> Recurring low-complexity workflows </div>
---	--

Protect the premium work. Test the repeatable work. Build the hybrid model.

Protect the premium work. Test the repeatable work. Build the hybrid model.

VALUE DECISION

THE PRICING LADDER

Pricing changes only after confidence increases. This is a ladder of confidence, not a philosophy.

The safest transition is not billable today, fixed fee tomorrow. It is a testing sequence.

**Hybrid pricing is not a belief system.
It is a testing sequence.**

BEFORE CHANGING PRICING

Efficiency only becomes margin when the workflow knows what to do with it. Before changing pricing, the firm should decide how saved time will be used.

Next: use the ladder to sequence risk and evidence.

FRAMEWORK

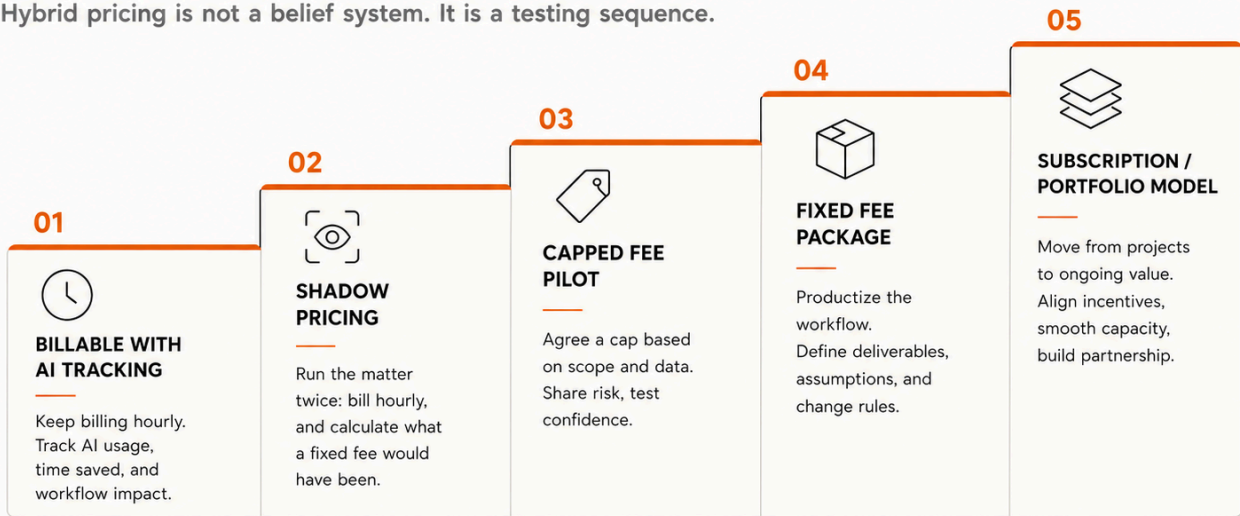
THE HYBRID PRICING LADDER

The ladder protects the firm from treating fixed fees as a leap of faith. Each step increases evidence about scope, variance, review load, exception rate, and client value.

PRICING EVOLUTION

THE HYBRID PRICING LADDER

Hybrid pricing is not a belief system. It is a testing sequence.



Flat fees are not about charging less.

They are about understanding the work clearly enough that efficiency can become margin without sacrificing trust.

PEOPLE

GOVERNANCE BEFORE SCALE

Governance is not legal theater. It is the infrastructure that decides what can be automated, what must be reviewed, what data is safe, who owns the output, and when the workflow must stop.

Scaling AI before governance is how small experiments become unmanaged risk. A governed pilot should define not only what the model does, but what the firm does around the model.

- Approved tools
- Data boundaries
- Review standard
- Escalation rules
- Audit trail
- Owner
- Stop condition

COMMON FAILURE MODES

- Sensitive data enters the wrong tool
- No one knows what AI touched
- Output is reused without review
- No escalation path exists
- No single owner is accountable
- The pilot expands before rules are stable

The question is not “Can the model do it?” The question is “Can the firm safely rely on the workflow around it?”

FRAMEWORK

THE TRAINING LAYER

AI should not make juniors and paralegals invisible. It should make the team more leveraged. That requires intentionally redesigning how learning happens when repetitive processing work gets compressed.

For a solo practitioner, AI may create leverage before the next hire. For a firm with a team, AI should make the team more capable, not simply smaller.

If AI compresses repetitive work, the firm needs to design new learning loops.

- Save before/after examples for common workflows
- Review AI outputs with juniors, not just partners
- Build issue-spotting libraries from real matters
- Document escalation decisions and exception patterns
- Turn corrections into reusable training material
- Move paralegals from processing to quality control and workflow ownership

The goal is not to remove training. The goal is to make training reviewed, reusable, and intentionally designed.

INTERNAL BENCHMARK

THE 30-DAY HYBRID AI PILOT









A pilot should create a decision memo, not just enthusiasm. At the end, stop, clean, expand or price-test.


AI VORTEX · HYBRID AI FIRM PLAYBOOK

EXECUTION ROADMAP

THE 30-DAY PILOT PLAN

A focused plan to test, learn, and decide with confidence.

 PHASE 1 DEFINE & SELECT DAYS 1-5	 PHASE 2 BUILD & PREPARE DAYS 6-15	 PHASE 3 RUN & MEASURE DAYS 16-25	 PHASE 4 DECIDE & SCALE DAYS 26-30
<p>OBJECTIVE Choose the right workflow and define success.</p> <p>KEY ACTIONS</p> <ul style="list-style-type: none"> • Run readiness scorecard • Identify a repeatable workflow • Define scope and boundaries • Set success metrics • Align stakeholders <p>OUTPUT  Pilot charter & success metrics</p>	<p>OBJECTIVE Build the workflow and prepare the team.</p> <p>KEY ACTIONS</p> <ul style="list-style-type: none"> • Map the workflow steps • Select and configure tools • Create templates / prompts • Set review and quality rules • Train the team <p>OUTPUT  Workflow ready for pilot</p>	<p>OBJECTIVE Run the pilot and capture real data.</p> <p>KEY ACTIONS</p> <ul style="list-style-type: none"> • Execute with real matters • Track time saved vs. baseline • Capture quality feedback • Monitor exceptions • Hold mid-pilot check-in <p>OUTPUT  Pilot data & insights</p>	<p>OBJECTIVE Decide the next step with evidence.</p> <p>KEY ACTIONS</p> <ul style="list-style-type: none"> • Review results vs. metrics • Refine workflow and rules • Decide pricing model • Document playbook • Plan the next workflow <p>OUTPUT  Go / No-Go decision & scale plan</p>



A 30-day pilot does not replace strategy.
It creates the evidence you need to choose your strategy.

The deliverable is a decision memo: continue, stop, clean the workflow, expand the pilot, or test pricing — based on variance, review load, and exception rate.

DECISION TOOL

METRICS THAT MATTER

The first benchmark that matters is not the market average. It is your own baseline. Measure before changing the business model.

If the only metric is hours saved, the firm will misunderstand the value.

Before changing pricing, capture the baseline:

- Average cycle time
- Average lawyer review time
- Number of handoffs
- Missing-input events
- Exception rate
- Escalation frequency
- Quality issues
- Client turnaround time
- Non-billable admin time
- Reusable knowledge captured

SPEED

- Cycle time
- Client turnaround

REVIEW

- Lawyer review time
- Escalation rate

QUALITY

- Error rate
- Quality issues

ECONOMICS

- Margin per matter
- Pricing variance

GOVERNANCE

- Compliance
- Auditability

TRAINING

- Reusable examples
- Judgment development

Do not price the promise. Price the measured workflow.

FRAMEWORK

PARTNER MEETING CHECKLIST

Use this before approving a pilot, buying a tool, changing pricing or rolling a workflow out to the wider team.

AI VORTEX · HYBRID AI FIRM PLAYBOOK

DECISION TOOL

BEFORE APPROVING AN AI PILOT, ANSWER THESE QUESTIONS

Use this before approving a pilot, buying a tool, changing pricing, or rolling a workflow out to the wider team.

1. Which workflow are we testing?
2. What exact input starts it?
3. What exact output should be produced?
4. What can AI do safely?
5. What must the lawyer review?
6. What cannot touch AI?
7. What triggers escalation?
8. What metrics are we tracking?
9. Are we changing pricing, shadowing it, or just measuring?
10. What would make us stop?
11. What would make us expand?



Most firms don't need more AI enthusiasm. They need a better operating sequence.

QUICK REFERENCE

QUICK REFERENCE SHEET

This page is designed to be screenshotted, printed or used in a partner meeting.

THE HYBRID SEQUENCE

- **Policy** → decide what can touch AI
- **Workflow** → map the work and the drag
- **Pilot** → test one bounded use case
- **Scorecard** → benchmark readiness and review load
- **Pricing** → adjust only after evidence
- **Governance** → standardize before scale

Stop if: data risk, quality risk, unclear scope, no reviewer, no owner.

Go if: scope, review, exceptions, governance, and client value are clear.

The goal: borrow AI-native operating logic without breaking pricing, training, governance, or trust.

NEXT STEP

CHOOSE THE FIRST WORKFLOW

This playbook gives you the operating sequence. The hard part is applying it to your firm's actual workflows, incentives, tools, people, and risk tolerance.

You do not need to become AI-native overnight.

You need to know which parts of the AI-native operating model are safe to borrow first, which parts should stay human-led, and which workflow is ready to test.

That is where a Hybrid AI Workflow Pilot starts.

HYBRID AI WORKFLOW PILOT

- Choose one high-fit workflow
- Map the current state
- Identify invisible work tax and capacity leakage
- Design the AI-assisted version
- Define governance, review and escalation rules
- Create an internal benchmark method
- Recommend whether to keep billable, shadow price, cap, fix or wait
- Deliver a partner-ready memo

The goal is not to buy another tool. The goal is to know what should be built, governed, tested, and measured before the firm scales AI around real client work.

If this helped reframe the AI conversation inside your firm, go back to the LinkedIn post and comment with the workflow you would map first.

Want help choosing the first workflow? Contact details on the next page.

CONTACT

WORK WITH MANU

If you want help selecting the first workflow, mapping the current state, and running a governed pilot that produces real evidence (not theater), reach out below.

CONTACT**Manu Ayala**

AI Architect for Law Firms

AI Vortex

Legal AI Infrastructure · Visibility · Workflow Intelligence

Website: aivortex.io/legal

LinkedIn: linkedin.com/in/aivortex

Calendly: calendly.com/manuel-aivortex/ai-infrastructure-workflow-audit

Email: manuel@aivortex.io

This playbook is designed to be used in internal partner discussions and workflow design meetings.