

# GenAI for Private Equity

JUNE 2026

# Contents

---

01	<b>Introduction</b>	3-4
02	<b>Vision on AI in PE</b>	5-9
03	<b>Past projects</b>	10-13
04	<b>Way of working</b>	14-16
05	<b>Vision on AI in PE (cont.)</b>	17-23

---

SECTION 01


# 01 Introduction

---

# Daniël Poolman

Founder, Poolman Ventures · Delft, NL



 **BSc Computer Science**

 **MSc Econometrics**

 **CLOUD SOLUTIONS**


 **VORTEX CAPITAL PARTNERS**

 **LEXAR**

**TODAY**

 **Poolman Ventures**

*Building (AI) solutions for private equity, grounded in engineering and deep knowledge of PE workflows.*



SECTION 02


# Vision on AI in PE

---

# The way AI is used within PE teams today and its problems

## THE LOOP

Works for simple tasks (summarize, draft, reformat). Stalls at 70-80% on anything complex.



## THE REAL PROBLEM

**You can't trust the output for real work.**

### Underlying problems

#### 1 Cold start every session

It cannot see Affinity, the data room, or last week's IC memo, and remembers none of your workflows, preferences, or corrections. Every session starts from zero.

#### 2 Hidden assumptions


It picks the structure, the weighting, and the criteria without checking. You catch the mismatch only in the final draft.

#### 3 Generic, not your edge

What comes back is the average of the internet: fluent, plausible, and the same answer any fund would get. It reflects none of your thesis or judgment.

# A worked example: Buy-and-build lists today

## THE WORKFLOW TODAY



### WHY AI CAN'T HELP HERE

## No place for an agent to plug in

Filtering is the only structured step, and it is the lowest-value one: the analyst already does it by hand. The judgment that matters, what "fit" means, lives in the associate's head.

#### **ASSOCIATE** Judgment locked in their head

What "fit" means is never written down. No artifact for an agent to match against, no checkpoint for the partner to validate.

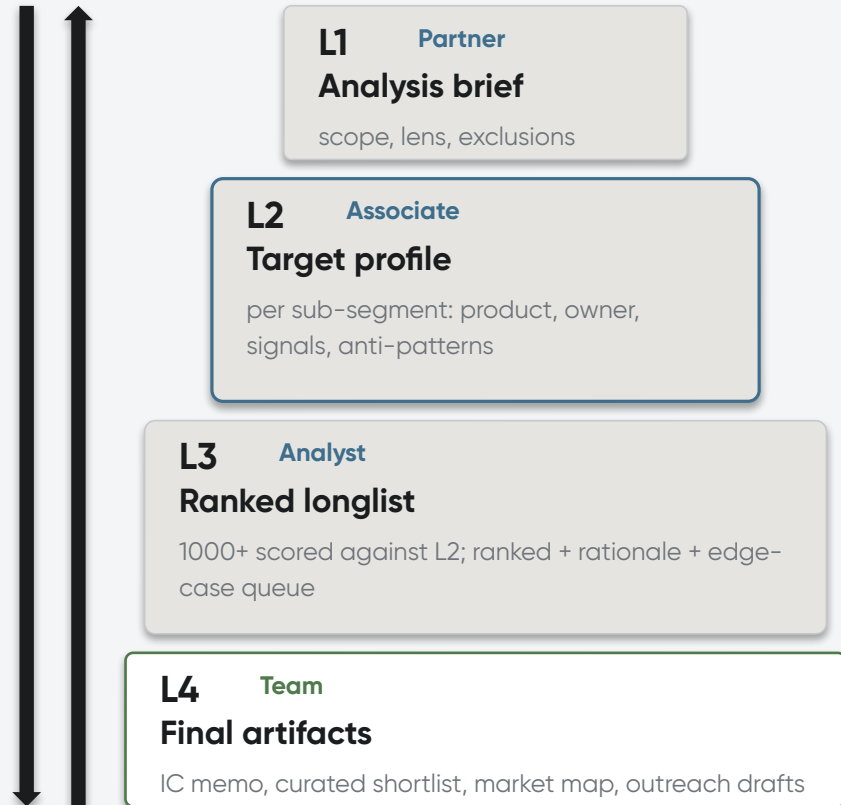
#### **ANALYST** Structured filtering only

Sector, geography and headcount filter cleanly. Product shape, founder profile and commercial readiness can only be eyeballed by hand, so 1000+ targets get a glance at best.

# A worked example: Buy-and-build lists redesigned

## FEEDBACK LOOPS

↓ refine + ↑ propagate (left): a layer edit cascades down, a finding propagates up  
 ○ within-layer (right): agent drafts, human reviews, agent revises



## WORKED EXAMPLE · buy-and-build for portco X

WITHIN LAYER

### ANALYSIS BRIEF · L1 artifact

<b>Goal</b>	3 platform bolt-ons for portco X within 18 months
<b>Deal lens</b>	bolt-on at €15-40m EV; minority OK for proof-of-thesis
<b>Scope</b>	vertical SaaS · NL / BeNeLux / DE · construction
<b>Exclusions</b>	pre-revenue · US-only · hardware-heavy · non-SaaS

### TARGET PROFILES · L2 artifacts · one per sub-segment

Field service mgmt	Estimating tools	Compliance / safety SaaS
<b>Product</b> dispatch + ops <b>Owner</b> founder, €2-8m <b>Signals</b> vertical, NL HQ	<b>Product</b> takeoff + quote <b>Owner</b> subscale ARR <b>Signals</b> BIM, NL/DE	<b>Product</b> safety + training <b>Owner</b> founder fatigue <b>Signals</b> regulated, SaaS

### RANKED LONGLIST · L3 artifact · scored against L2 target profiles

Company	Segment	Match score	Rationale
FieldServiceCo NL	Field service	92	exact fit · NL-HQ · founder-led estimating · BIM ties · subscale safety SaaS · founder fatigue adjacent vertical · off-thesis ranked vs L2 target profiles
EstimaXYZ BE	Estimating	8.4	
ComplyHub DE	Compliance	7.8	
DispatchPro NL	Field service	7.1	
... 1000+ rows			

### FINAL ARTIFACTS · L4 · downstream deliverables

<b>IC memo</b> deal-specific writeup	<b>Curated shortlist</b> 10-20 for partner review	<b>Market map</b> by sub-segment	<b>Outreach drafts</b> per-target email + brief
---	--	-------------------------------------	--

# Discussion

## WHERE YOU ARE TODAY

- What are your biggest operational challenges?

---

- What are you already doing with AI?

---

- Where does your data live, and how structured is it?

## WHERE AI COULD PAY OFF

- Where do you see the highest-ROI use cases?

---

- Which workflow would you want to tackle first?

---

- What would make a first pilot a success?

SECTION 03

# Past projects

---

# A production value-creation platform

## MODULE 01

### Investment management overview

- Valuations engine + bridge-import
- Quarterly portco reporting and roll-up
- Legal-structure graph: entities, directors, KYC

## MODULE 02

### Portfolio collaboration

- Playbook-driven task management (Kanban)
- Document sharing (SharePoint integration)
- Contacts and teams: Lexar staff + portco operators

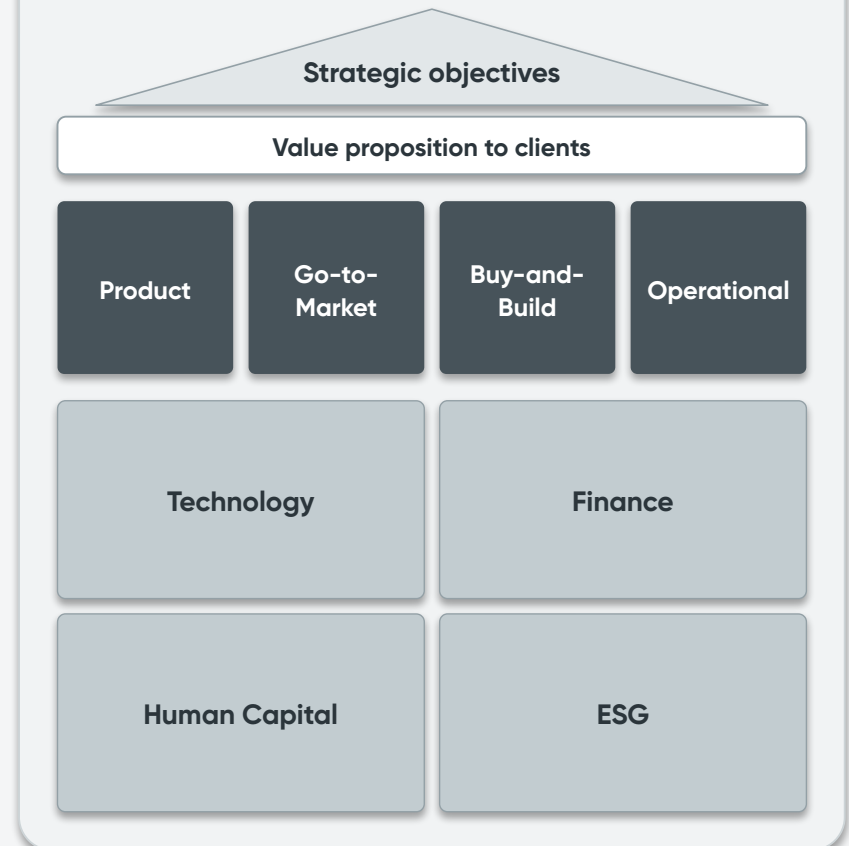
## MODULE 03

### Value creation

- 8-pillar maturity tracking (5-level scale)
- QRAP portco questionnaires (incl. ESG)
- Playbooks: project to task to subtask scaffolds

## VCP FRAMEWORK

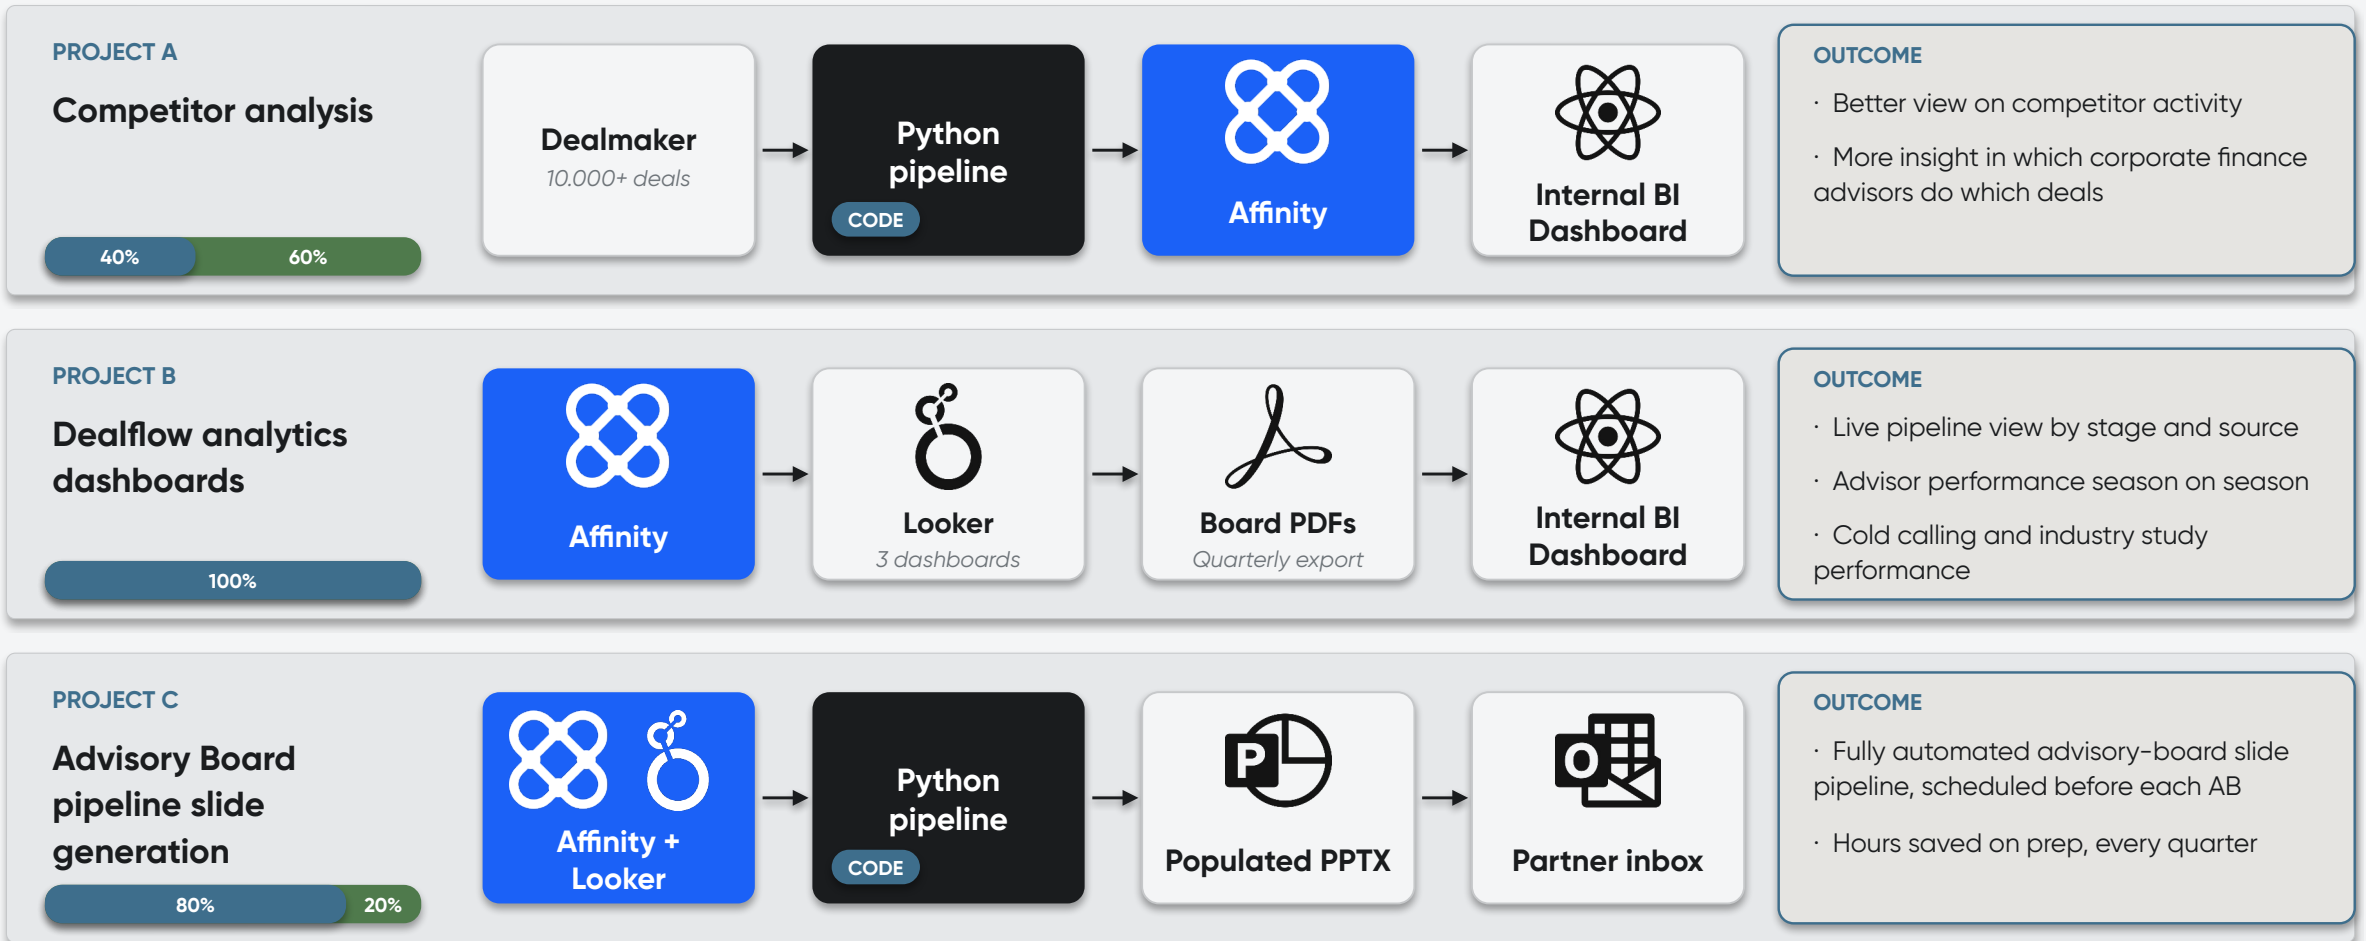
Client impact  Enabler



**LIVE**

A cockpit for both investment teams and portfolio companies to centralize important workflows

# Three reporting and analysis workflows automated end-to-end



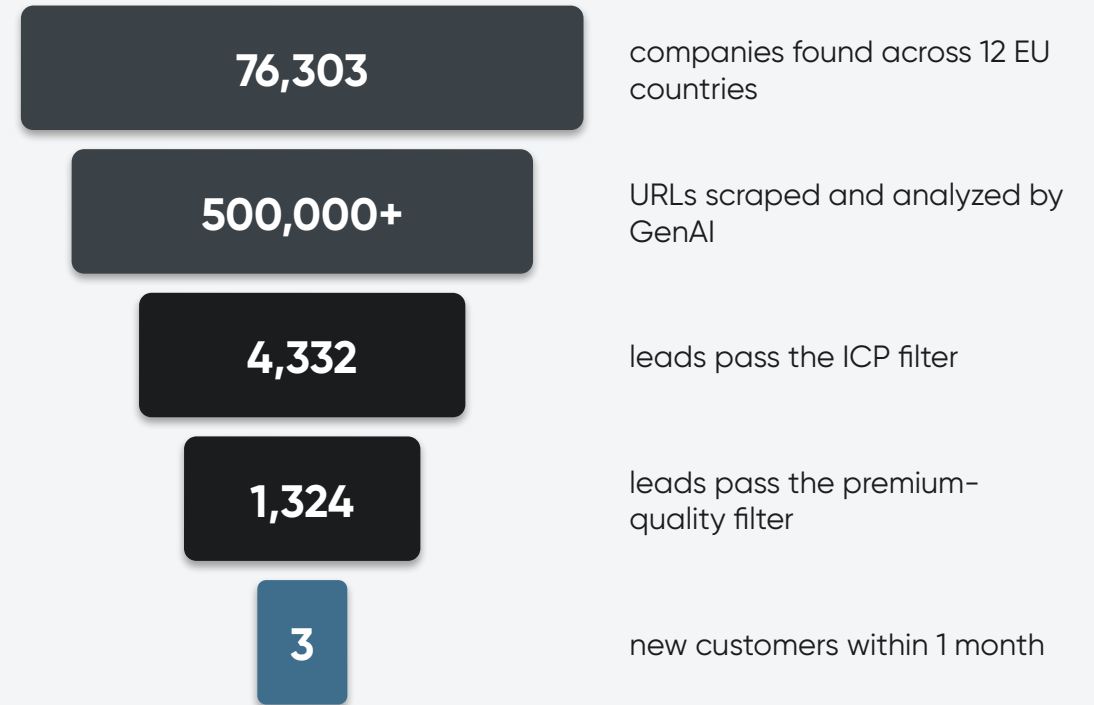
DATA AI / AGENTIC

# Manual lead sourcing replaced: how TMM mapped the EU premium-meat market

## AUTOMATED PIPELINE



## RESULTS



## OUTCOMES

Sales focuses on closing

Org-wide lead-gen standard

Complete addressable market in EU mapped

SECTION 04

# 04 Way of working

---

# Way of working



# Recurring meetings during projects to keep aligned on vision, scope, budget and output

## DURING ENGAGEMENT

### Working session

**Weekly** 60–90 min

FIRM-SIDE

**Analyst + Investment Manager**

DECIDES

*Scope on the table, design, build, blockers.*

### Steering committee

**Biweekly** 30–45 min

FIRM-SIDE

**Partner + working-session group**

DECIDES

*Roll-up of decisions, scope, risk, budget.*

### Milestone review

**At gates** 60 min

FIRM-SIDE

**Full sponsor group + IC member**

DECIDES

*Kickoff, mid-engagement, pre-handoff sign-off.*

ONCE WORKFLOWS ARE LIVE

## Maintenance

### HOSTING

*Cloud infra running, uptime monitored*

### BUG FIXES

*Patches when a workflow misbehaves*

### MODEL ROTATIONS

*Provider + model swaps handled*

### EVAL MONITORING

*Drift and accuracy alerts*

### COST GUARDRAILS

*Token spend stays in budget*

SECTION 05

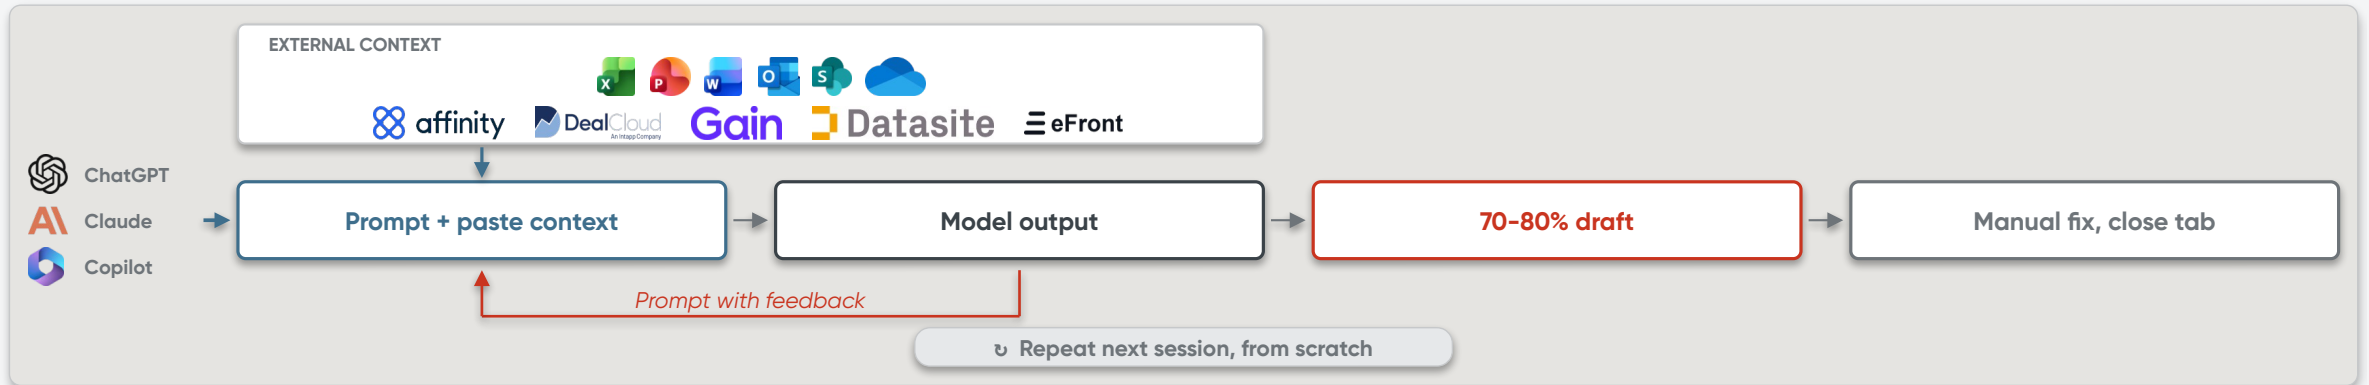
# Vision on AI in PE (cont.)

---

# The way AI is used within PE teams today and its problems

## THE LOOP

Works for simple tasks (summarize, draft, reformat). Stalls at 70-80% on anything complex.



## THE REAL PROBLEM

**You can't trust the output for real work.**

### Underlying problems

#### 1 Cold start every session

It cannot see Affinity, the data room, or last week's IC memo, and remembers none of your workflows, preferences, or corrections. Every session starts from zero.

#### 2 Hidden assumptions

It picks the structure, the weighting, and the criteria without checking. You catch the mismatch only in the final draft.


#### 3 Generic, not your edge

What comes back is the average of the internet: fluent, plausible, and the same answer any fund would get. It reflects none of your thesis or judgment.


# Root cause 1: we treat AI like a tool, not like a team member

An LLM is a brilliant new hire on day one. Sharp, well-read, fast. What it lacks is everything a person absorbs over years inside a firm. We onboard every human; we give AI a prompt, then expect senior work.

## A NEW HIRE



## AI TODAY



WHY YOU SAW THOSE SYMPTOMS · [back to slide 06](#)

**2** **Assumptions you never see**  
No context supplied, so the model fills the gaps silently.

**3** **The conversation drifts**  
No ways of working to anchor the goal as it runs.

# Root cause 2: PE's digital ops are not structured for automation or AI

## 1 Systems are siloed

Each SoR holds a fragment of the deal. Affinity doesn't see the data room; the tracker doesn't link to IC memos. No system holds the whole picture.

## 2 Knowledge is implicit

Which list is the live pipeline, what "Stage 3" means, which template fits: the meaning lives in people's heads, never encoded where a system can use it.

## 3 Work is manual

The systems offer no way to package a workflow, trigger it, or carry its output forward. So every run is driven by hand, and forgotten.

### PRODUCTIVITY LAYER

MICROSOFT 365		NL DATA	C	G	
GP ACCOUNTING		E-SIGNATURE			

### THIN BY DESIGN

Project management		Knowledge management	
--------------------	--	----------------------	--

### SYSTEMS OF RECORD

CRM / relationships	Fund reporting / accounting	Deal sourcing	Data rooms / VDR

# Closing both gaps requires a new way of working and a new layer of infrastructure

## DESIGN FLAW

No structure in how we work with AI

### 1 Break workflows down into reviewable layers

Decompose workflows into an abstraction hierarchy. Each layer produces a distinct artifact. Each artifact gets reviewed before the next layer proceeds.

WITHOUT THIS

*Clean process, still pasting context into a chat window.*

Slide 10 unpacks this

## OPS GAP

Systems not structured for agents

### 2 Make context and workflows persistent and reusable

A layer on top of your existing ops stack that holds persistent context (the data every run reads and writes) and packages reusable workflows (the runs themselves), for both humans and agents.

WITHOUT THIS

*All systems connected, still reviewing blind at the end.*

Slide 11 unpacks this

## COMPOUNDING SYSTEM

Both are required. Together they create a compounding system: each run produces reviewed, reusable artifacts that make the next run faster, more accurate, and more aligned.

# Break workflows down into reviewable layers...

Layer	Artifact	Seniority
L1	<b>Narrative</b> Why we're doing this, for whom, why now	<i>Partner-level decisions</i>
L2	<b>Spec</b> Evaluation criteria, structural choices, output scope	<i>Partner + associate shaping</i>
L3	<b>Wireframe</b> Structure, key claims, data points	<i>Associate + analyst building</i>
L4	<b>Final artifact</b> The memo, deck, scored longlist, drafted email	<i>Analyst-level execution</i>

## WHAT THIS DOES

### 1 Surfaces assumptions

Left alone, the model silently picks the framing, the criteria, the depth. Layering turns those into artifacts you approve at L1 and L2, before anything is built on them.

### 2 Makes each task tractable

"Do a full industry study" is too big to steer or check. Split into layers (brief, then specs, then scored list), each step is small enough to get right and to verify.

### 3 Creates feedback architecture

Not a one-way pipeline. A correction or new signal enters at the layer it belongs to and re-flows from there, so you fix the cause, not rerun the whole job.

# ...and structure context and workflows such that they are reusable

## REUSABLE WORKFLOWS

### Skills


The procedure for a recurring task: the steps to source, scope a CDD, or draft an IC memo.

### Triggers

The event or schedule that starts a run (a deal hits Stage 3, a quarter closes), no prompt needed.

### Guardrails

The scoped permissions and sign-off gates that bound each run: a skill touches only what its task needs.



## PERSISTENT CONTEXT

### Records


Deal, portco and thesis facts, assembled from your SoRs into one queryable place the agent reads and writes.

### Conventions

The definitions and criteria the agent looks up: what "Stage 3" means, the ICP criteria, which template, house style.

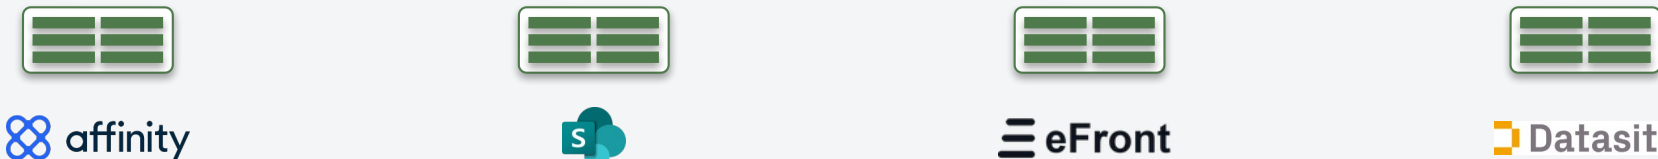
### Artifacts

The versioned outputs of past runs: briefs, ICP specs, scored longlists and corrections.



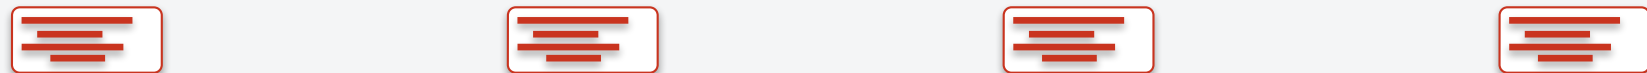
## AGENTIC

structured & connected



## TODAY

raw, unstructured, siloed





**Poolman Ventures**