


Agentic Consulting Framework

Deciding what to point agents at — before anyone writes a line of code



The Conversation Happening Across Enterprises Right Now

Every enterprise IT and business leader has been told they need agents. Most have run a pilot. A meaningful share of those pilots are now stuck in the same place: an impressive demo, a happy path that held, and a rollout that stalled the first time an input arrived in an unexpected format, a system turned out to have no callable API, or a risk committee asked who signs for a machine decision.

⚠ That failure is rarely technical. It is a **selection failure**. Nobody measured whether the process could carry autonomy before the autonomy was built.

This failure costs organizations in three ways:

Blueprint Debt

Blueprint gets fed whatever documentation is nearest, and faithfully reproduces legacy technical debt as a modern case type — faster than a human could have made the same mistake.

Integration Debt

Integration debt surfaces in build rather than in the roadmap, where it could have been priced and sequenced correctly.

Governance Failure

The pilot that stalls in governance is the pilot stakeholders remember at renewal or budget planning — poisoning the well for future investment.

We fix the selection problem. Not with a workshop and a deck — with a scored, reproducible instrument that names which processes should become agentic on Day 1, which need a year of enablement first, and which should never be agentic at all.

The Idea in One Paragraph

Deterministic transformation asks how to eliminate variance. Agentic orchestration inverts that: an agent earns its cost only where variance cannot be engineered away — unpredictable inputs, meaning that must be interpreted, paths decided case by case, exceptions as the norm rather than the tail.

A process a BPM engine already handles cleanly is, almost by definition, a poor agentic candidate. So the instrument has to measure two things at once, and hold them apart:

Agentic Need

How much irreducible cognitive work the process carries. Measured across:

- Cognitive variability
- Data semantic richness
- Exception density

Agentic Feasibility

Whether the surrounding systems, data and controls can actually support an autonomous actor. Measured by:

- Digital tool and API maturity
- Organizational risk appetite

Cross those two axes and the roadmap draws itself.

What We Measure

Four dimensions, each decomposed into the same four categories — which is what makes them comparable against one another rather than four unrelated scores.

Dimension	What It Asks	Axis
Cognitive & variability	How much the process depends on human cognition, and how variable its inputs, decisions and paths are.	Need
Data semantic richness	How much meaning, context and interpretation the data carries.	Need
Digital tool & API maturity	How accessible and automatable the supporting systems are.	Feasibility
Exception density	How often the process leaves its happy path.	Need

The Instrument: 56 Behavioural Statements

Fifty-six behavioural statements underneath, weighted by the enterprise's own priorities through a pairwise comparison of the four dimensions and a percentage split within each. The statement set is configurable per industry and per engagement.

- ☐ The statements are the reason this survives a sceptical COO or Chief Architect. None of them asks "are you ready for AI?"

Each statement describes an **observable operational fact** the respondent can verify from their own week:

Input Variance Example

"The inputs triggering this process arrive in highly unpredictable formats rather than standardised forms."

Decision Complexity Example

"An operator must synthesise information from multiple disparate documents, policies or past cases before rendering a decision."

Perceptions can be argued with. Observations can be triangulated against process mining. This distinction is what makes the instrument defensible in a governance conversation.

What the Enterprise Experiences

The Experience

Ten minutes. No login. A branded microsite stakeholders can complete before a first meeting — which means the discovery conversation starts with data instead of with a blank page.

Immediate In-Browser Output


- A weighted readiness index
- A heatmap sorted weakest-first
- A radar overlaying readiness against strategic importance
- A category-level breakdown an architect can build from



10-Minute
Assessment

Instant Score
& Heatmap

Data-Driven
Discovery



What Comes Out: A Live Run

LIVE EXAMPLE OUTPUT

69

Overall Score

"Solid, scaling with focus."
Composite agentic candidacy
index out of 100.

3.44

Overall Maturity


Out of 5.0 across all four
dimensions in this live run.

3.1

Lowest Score

Digital tool & API maturity — the
sharpest pain point identified.

Dimension	Score	Weight	Readiness
Digital tool & API maturity	3.1 / 5	7	63%
Cognitive & variability	3.4 / 5	12	69%
Exception density	3.4 / 5	11	69%
Data semantic richness	3.9 / 5	6	77%



Reading the Scores Correctly

Two different answers, and cross-functional teams need both.

By Raw Score: The Technical Constraint


The heatmap names **API maturity** as the sharpest pain point, because it has the lowest raw score at 3.1 / 5. This is the binding technical constraint that must be resolved before deployment.

By Gap (Importance – Readiness): The Value Target

Cognitive & variability leads at 31 points gap. Exception density is second at 23 points. This is where the business value sits. Fix the constraint, aim at the value.

- A high score on cognitive variability, semantic richness or exception density describes a process that is **hard for deterministic automation** – which is the strongest possible signal of agentic need, not a deficiency. Read the composite as an **agentic candidacy index**.

A leader who hears "69/100 readiness" argues with the grade. A leader who hears "your processes are strongly agentic-shaped and your integration layer is the constraint" engages with the finding.



Where the Roadmap Comes From

The Need × Feasibility matrix produces four distinct strategic recommendations. Each quadrant has a different answer — and the wrong quadrant is where most failed pilots live.

	Low Feasibility	High Feasibility
High Need	Year 1 — invest to enable. The value is real; the plumbing is not. Integration and data work first, agent second. This quadrant is the classic failed-pilot trap.	Day 1 — flagship missions. Highest value, lowest risk of technical failure. The first supervised mission goes here.
Low Need	Neither. Do not spend agentic budget here.	Deterministic automation. Classic workflow. An agent adds cost and non-determinism for no gain.

Day 1 vs. Year 2 Mission Design

The worked example lands in high need with constrained feasibility. The correct recommendation is not "deploy agents" — it is to close the integration gap on the two processes that matter, then deploy on those two. A bigger, better-sequenced programme than the pilot the organization was about to buy.

	Day 1 Mission	Year 2 Mission
Autonomy	Assisted. Agent proposes, human disposes.	Delegated within an envelope; human handles escalations.
Selection Rule	High need, high feasibility, bounded blast radius, reversible actions.	Feasibility created during Year 1; irreversible or regulated actions only under proven guardrails.
Guardrails	Deterministic case lifecycle, mandatory approval gate, full decision logging.	Policy-based autonomy envelope, exception-triggered hand-back, continuous evaluation.
Measured By	Decision quality vs human baseline; containment rate.	Straight-through rate; cost per transaction; escalation trend; audit findings.

Harmonizing with Your Pega Ecosystem

FOR PEGA EXPERTS & ENTERPRISE ARCHITECTS

We are not a competing platform. We do not compete with anything in the Pega stack. We provide the discovery instrument **in front of Pega GenAI™ Blueprint** and the guardrail design **behind** it. Six specific effects on an enterprise architecture and joint roadmap:

Blueprint Precision

Curated seed packs instead of whatever documentation was nearest. Case models that encode business priority, and materially less rework after generation.

A Defensible Autonomy Story

Risk, compliance, and enterprise architects get a quantified basis for every autonomy decision. That is the difference between a pilot that passes governance and one that dies in it.

Integration Spend in the Right Year

API maturity is measured before the agentic commitment, not discovered during build.

Fewer Failed Pilots

The low-need / high-feasibility quadrant is explicitly excluded, stopping the most common category of wasted agentic spend.

Expansion Mapping

Scoring across a function shows where the next three missions are. One case type becomes a scalable programme.

A Re-baselining Instrument

Re-run it after each phase and you have a clear movement narrative — empowering transparent ROI discussions and subsequent expansion.

How It Lands on the Platform

For the architects in the room, our design position rests on five statements:

The Case Lifecycle Is the Guardrail

The agent is free inside a stage; stage transitions, approvals, SLAs and case history are not negotiable. That is what makes autonomy defensible to a risk committee, and it is the architectural reason Pega hosts agentic work rather than bolting it on.

We Build with Infinity '26, Not Around It

Pega as an MCP server is exactly the pattern our reference architecture calls for — tools exposed for discovery and execution, with the agent treated as an untrusted requester. Agentic Process Fabric's registry, with its per-workflow discoverability flag, is the native exposure control we leverage.

A Deliberate Position on Predictable AI

Pega front-loads reasoning to design time — right for cost and determinism. Our default architecture keeps cognition outside Pega, reached over MCP, so the case stays deterministic and the economics remain predictable. Where scores show low variability, we explicitly recommend the Pega-native path with no external component.

01

Every Effectful Connector Gets a Deterministic Gate

A when rule evaluated before any state-mutating call, outside the reasoning path, reading the enterprise's own constraint set. No effectful connector is reachable except through it.

02


Tools Exposed in Three Waves, by Effect

Read-only first at zero blast radius, then staged actions under a soft checkpoint, then state-mutating calls behind an approval gate with an idempotency key. A Wave 1 deployment delivers value in about six weeks while mutating nothing.

Category Signal from the Assessment

Each category signal from the scoring instrument maps directly to a specific architectural requirement. The assessment output is not a report — it is a design brief.

Category Signal	What the Architecture Must Then Provide
High input variance	Intake classification, confidence thresholds, reject-to-human below threshold.
High decision complexity	Retrieval over an authoritative corpus, rationale capture, deterministic policy checks on top of model output.
High pathway fluidity	Legitimate alternate stages; every agent-chosen path expressible as a case state, never an untracked side effect.
High exception density	Explicit exception taxonomy; per-exception disposition; hand-back semantics.
Low API maturity	Integration remediation first. Robotics only as a time-boxed bridge with a decommission date.



The Assessment-to-Blueprint Handshake

A seven-step process that converts scored assessment output into a production-ready Blueprint seed pack and guardrail overlay.


#	Step	Output
1	Rank candidate processes by Need x Feasibility	Ranked candidate list
2	Cut at the feasibility gate	Day 1 shortlist + Year 1 enablement list
3	Assemble a curated seed pack per Day 1 process	SOPs, exception logs, policy documents, integration inventory
4	Generate the workflow model in Pega GenAI Blueprint	Draft case type, stages, steps, data model
5	Apply the guardrail overlay from the category scores	Approval gates, exception taxonomy, autonomy envelope
6	Mark the agentic execution zones	Which stages are agent-executed, which stay deterministic
7	Define the evaluation harness before build	Acceptance criteria and production monitoring plan

- ☐ **Step 3 is the whole argument.** Blueprint is very good at faithfully reproducing what it is given. The assessment is what earns a document its place in the seed pack.

The Engagement Model

Five stages from self-service assessment to supervised autonomy in production. The assessment is free. The first commercial decision is which single mission to build – a far easier decision to get funded than an ambiguous "agentic strategy".

Stage	Duration	What the Enterprise Gets
Assessment	10 minutes, self-service	Weighted readiness index, heatmap, radar, category detail. No cost, no login.
Multi-stakeholder read-out	1 session	The same process scored by three roles. Where leadership and the front line disagree is usually worth more than the score.
Ambition Lab	1–2 days	Signed Day 1 / Year 2 mission map, autonomy envelope per mission, dependency register, and an explicit not agentic list.
Blueprint sprint	2–4 weeks	Curated seed pack, generated case design, guardrail overlay, evaluation harness specification.
Day 1 mission build	6–12 weeks	Supervised autonomy in production, measured from the first commit.



Questions We Ask

USEFUL IN ANY ORGANIZATIONAL CONTEXT, WHETHER OR NOT WE ARE IN THE ROOM

These five questions are diagnostic instruments. An organization that can answer all five with evidence is ready to build. An organization that cannot is not — and knowing that before the build is the entire value of the framework.

1

Straight-Through Rate

What percentage of your cases go straight through today — and can you evidence it?

2

Reversibility

Which of these systems can an automated actor undo an action in?

3

Audit Trail

When your agent makes a decision, where does the rationale get written, and who retrieves it a year later in an audit?

4

Explicit Exclusions

Which of your processes should not be agentic, and how did you decide?


5

Improvement Signal

How will you know the agent is getting better, rather than just getting used more?

How We Keep It Honest

Three things we say to our partners and clients before they ask, because a framework that hides them does not survive a technical buyer.



It Is Self-Reported

Scores are perceptions until corroborated. We triangulate against process mining, straight-through rates or a sample audit before any Day 1 mission is committed.

Scope It to One Named Process

An organization-level score is a conversation starter. A process-level score is a decision.

The Regulatory Clock Moved; the Work Did Not

The EU AI Act's high-risk obligations were deferred by the Digital Omnibus to December 2027 and August 2028. That removes the deadline, not the requirement — and conformity evidence for human oversight and quality management is longitudinal. No one can manufacture fifteen months of behavioural evidence in the final quarter of 2027. **The deferral is the build window.**

What We Already Have

Not a concept. A working system at build depth across every layer of the framework.



Working Scoring Engine

A microsite with a live scoring engine. No login required. Immediate output in-browser.



100-Topic Architecture Framework

A complete architecture framework covering the full breadth of agentic enterprise design decisions.



Pega Capability Mapping

Complete mapping against Infinity '26 with honest Native / Partial / Gap verdicts – no marketing language.



Tiered Tool Exposure Playbook

Three-wave exposure model by effect type, with reference architecture and deployment sequencing.



Ambition Lab Agenda

A structured 1–2 day facilitation agenda producing a signed mission map and dependency register.




Two Product Designs at Build Depth

Designs for the gaps Pega does not currently cover – built to specification, not to concept.

The Agentic Candidacy Index at a Glance

A summary of the scoring logic: three dimensions drive need, one drives feasibility, and the quadrant intersection determines the recommendation. The index is not a readiness grade – it is a candidacy signal.



The Regulatory Context

The EU AI Act's high-risk obligations create a longitudinal compliance requirement that cannot be manufactured retroactively. The deferral timeline makes the current window the build window — not a reason to wait.

Key Dates

- **December 2027:** First deferred deadline under Digital Omnibus
- **August 2028:** Second deferred deadline
- **Now:** The window to build longitudinal behavioural evidence

What This Means Operationally

Conformity evidence for human oversight and quality management is longitudinal. No organization can manufacture fifteen months of behavioural evidence in the final quarter of 2027.

The deferral removes the deadline. It does not remove the requirement. Organizations that begin building governed agentic systems now will have the evidence trail. Those that wait will not.

⚠ The deferral is the build window. Treat it as one.

The Next Step



One Core Process. Ninety Minutes.

Pick a use case where an agentic conversation has started and stalled. We will run the assessment with cross-functional stakeholders, present the deltas, and produce the Day 1 / Year 2 mission map.

What Happens

Run the assessment with cross-functional stakeholders and present the scored deltas across roles.

What You Get

A signed Day 1 / Year 2 mission map with an explicit not-agentic list — and a clear sequencing rationale.

Our Commitment

If the scores say the process should stay deterministic, we will say that — saving both the enterprise and the implementation team from a pilot destined to struggle.