DEALCRAFT ADVISORY

THE AI-FIRST ERP EVIDENCE PLAYBOOK

AI-first ERP: a real delivery model, or a slide?

How buyers and bidding partners can test an AI-first claim in the plan, the price and the contract. A phase map, a savings worksheet, a scorecard and a clause checklist.

For CFOs, CIOs, procurement teams and programme sponsors evaluating ERP proposals, and for the SAP partners writing them. Print it, score two bids side by side, take it into clarifications.

dealcraftadvisory.ie

EXECUTIVE SUMMARY

The tools are real. Now test the claims.

Agents are now built into the SAP transformation toolchain. Within months, every ERP proposal will call itself AI-first.

In September 2026 SAP announced agents across its Business Transformation Management portfolio: a Process Consulting Agent in Signavio that mines processes across SAP and non-SAP systems, an AI Agent Hub mapping agents, models and MCP servers to capabilities and owners, LeanIX assistants for architecture and EU AI Act risk classification, UI-native agents in WalkMe, and Cloud ALM assistants across analysis, configuration, testing and rollout that also track agent actions in production.

>40%

of agentic AI projects predicted to be cancelled by end-2027

Gartner, June 2025

~130

of thousands of "agentic AI" vendors are the real thing

Gartner, June 2025

7

assistants in SAP Cloud ALM's agent-led migration toolchain

SAP, June 2026

57%

say data migration regularly delays their SAP projects

Xmateria via SAPinsider, 2026

Gartner gives three reasons for the cancellations it expects: escalating costs, unclear business value and inadequate risk controls. Each one can be tested in a proposal before signature.

FIVE THINGS TO TAKE AWAY

1. AI-first is a difference you can find. In the plan, the price and the contract, not in the executive summary.

2. Follow the saving. If agents remove a quarter of the effort, the price, the declared licences or the explanation should account for it.

3. Every agent output needs an owner. A named human reviewer on each side, before it becomes a deliverable.

4. Govern agents from sprint one. A register of what each agent does and which data it reaches, plus clear EU AI Act responsibility.

5. Some things don't change. Data runs still end with a reconciliation the business signs. Agents speed up the work around it.

TOOL 1 · THE PHASE MAP

Seven places AI should show up.

Tooling reflects SAP's 2026 announcements. Use the middle columns in clarifications; use the last column in the contract.

PHASE	TOOLING IN PLAY	WHAT SHOULD CHANGE IN THE PLAN	EVIDENCE TO ASK FOR	OWNER OF THE OUTPUT
1 · Discover	Signavio Process Consulting Agent: process mining across SAP and non-SAP systems	Shorter as-is analysis; fit-to-standard workshops start from mined processes	Processes to be mined, system access needed, time to findings	Business process owners validate
2 · Architect	SAP AI Agent Hub; LeanIX AI assistants incl. EU AI Act risk classification	Agent register and AI governance workstream from sprint one	Who maintains the register and classifies each AI use	Client architecture / CIO office
3 · Build	Cloud ALM assistants: system analysis, custom code, configuration, process	Fewer days for code analysis and configuration documentation	Share of agent-assisted remediation; human review rate	Partner solution architect signs off
4 · Data	Cloud ALM data management assistant	Faster profiling and mapping; same number of full data runs	How agents are used; unchanged load and reconciliation targets	Business data owners sign reconciliations
5 · Test	Cloud ALM AI-powered testing and clean core management	Earlier, faster test design; regression repeatable each release	Share of generated tests; review method; audit trail	Client test manager (and Quality, if regulated)
6 · Adopt	WalkMe UI-native agents and in-app guidance	Training shifts to in-app guidance; adoption measured	Adoption metrics; who pays licences after go-live	Business change leads; cost in business case
7 · Run	Cloud ALM tracking agent actions in production	Application management includes agent monitoring	Who watches deviations; service level for acting	Provider monitors; process owners decide

From the field. In most proposals we review, AI appears in three places: adjectives in the executive summary, an accelerators appendix and a slide of logos. It rarely appears in the effort model. Start reading there.

TOOL 2 · SAVINGS WORKSHEET

Where did the AI saving go?

The simplest test, and the one most often missed. Some of a claimed saving may fairly go to tool licences, a stronger team or extra assurance. None of it should simply disappear.

WORKSHEET		BIDDER A	BIDDER B
A	Baseline effort without agents (person-days)		
B	Blended day rate (€)		
C = A × B	Baseline cost		
D	Effort reduction claimed (%)		
E = C × D	Value of the effort agents remove		
F	Price reduction visible against the baseline (€)		
G	AI tool licences declared in the bid (€)		
H = E − F − G	Unexplained difference: ask where it went		

WORKED EXAMPLE

6,000 person-days at €900 is a €5.4M baseline. A claimed 25% effort reduction is worth €1.35M. The visible price reduction is 5% (€270,000) and declared licences are €250,000, leaving €830,000, about 61% of the claimed saving, unexplained. That may be margin, contingency or extra assurance, all legitimate, but the bidder should show where, phase by phase.

Questions for clarifications

- [ ] Show the effort model with and without agents, phase by phase.
- [ ] Which licences are included, for how long, and who holds them after go-live?
- [ ] If the claimed saving is not achieved, who carries the extra effort?
- [ ] If it is exceeded, how is the benefit shared?

TOOL 3 · PROPOSAL SCORECARD

AI-first, AI-assisted, or AI bolted on?

Score each question 2 = yes, 1 = partly, 0 = no. Use it as an evaluation sub-criterion, or, if you are bidding, as a red-team checklist.

#	QUESTION	BIDDER A	BIDDER B
1	**The plan is shorter or leaner in the phases where agents are used.**	0 1 2	0 1 2
2	**The price shows the effort reduction, or explains where it went.**	0 1 2	0 1 2
3	**Each agent and tool is named, with licence ownership during and after the programme.**	0 1 2	0 1 2
4	**Every agent output that becomes a deliverable has a named human reviewer.**	0 1 2	0 1 2
5	**There is an agent register: what each agent does, which data it reaches, who owns it.**	0 1 2	0 1 2
6	**Responsibility for EU AI Act classification and documentation is assigned.**	0 1 2	0 1 2
7	**Savings claims come with evidence from a comparable programme, not a demo.**	0 1 2	0 1 2
8	**A baseline is agreed so the savings can be measured.**	0 1 2	0 1 2
9	**There is a no-cost fallback if an agent underperforms.**	0 1 2	0 1 2
10	**Data used by agents stays within agreed residency and security rules.**	0 1 2	0 1 2
	Total out of 20		

SCORE	READING
16–20 · AI-first	The claim is in the plan, the price and the contract. Check the comparable evidence before relying on it.
9–15 · AI-assisted	Real use of AI, but parts of the claim are not yet in the plan, price or contract. Raise the gaps in clarifications.
0–8 · Bolted on	The AI story lives in the narrative. Score the conventional plan and ask for evidence on the rest.

TOOL 4 · CONTRACT AND BID CHECKLISTS

Settle it before signature.

Most of the risk behind Gartner's cancellation forecast, unclear value and weak risk controls, is cheaper to settle in the contract than after go-live.

Contract clause checklist (buyers)

- ☐ **Named tools and licence ownership:** which agents and platforms are used, who pays during the programme and after it.
- ☐ **Human sign-off:** every agent output that becomes a deliverable has a named reviewer on each side.
- ☐ **Agent register:** what each agent does, which data it reaches and who owns it, kept current from sprint one.
- ☐ **Regulatory responsibility:** who classifies AI use under the EU AI Act and keeps the documentation.
- ☐ **Measured baseline:** the effort and timeline savings are measured against, and how savings beyond plan are shared.
- ☐ **Fallback:** if an agent underperforms, the work is done conventionally at no extra cost.
- ☐ **Data protection:** agent access to personal and production data within agreed residency and security rules.
- ☐ **Exit:** agent configurations, prompts and the register are handed over at the end of the contract.

AI-first bid checklist (partners)

- ☐ Effort model shown with and without agents, phase by phase.
- ☐ Licences priced openly, with ownership after go-live.
- ☐ Named reviewers for agent output, on your side and the client's.
- ☐ One comparable programme with measured results, not a demo.
- ☐ Governance: agent register, EU AI Act approach, monitoring in run.
- ☐ A clear statement of what agents do not change, such as data reconciliation sign-off.

For partners. Evaluators facing five proposals that all say "AI-first" will score the one that proves it. Evidence beats volume.

SOURCES

Where the facts come from.

1	**Futurum Group**, SAP Links AI Agents to Business Processes and Oversight, on SAP's announcements at the Transformation Excellence Summit, September 2026. Customer time savings quoted there are not independently verified. futurumgroup.com/?p=95960
2	**SAP News**, Navigating the Transition from SAP Solution Manager to SAP Cloud ALM, June 2026. news.sap.com/?p=243739
3	**Gartner**, prediction that over 40% of agentic AI projects will be cancelled by the end of 2027, and the agent-washing estimate, June 2025, as reported by W.Media.
4	**SAPinsider**, reporting Xmateria's research "The True Cost of Data Procrastination", July 2026.

ABOUT THIS PLAYBOOK

The phase map, scorecard thresholds, worksheet and checklists reflect our practitioner experience across SAP and ERP programmes, not a published standard. SAP, Signavio, LeanIX, WalkMe and Cloud ALM are trademarks of SAP SE or its affiliates. Dealcraft Advisory is not affiliated with SAP or any organisation cited.

Related reading

The interactive version of these tools: dealcraftadvisory.ie/insights/ai-first-erp-delivery-model

How Many Data Runs Does an ERP Implementation Really Need? dealcraftadvisory.ie/insights/how-many-data-runs-erp-implementation

Starting Your ERP Journey: Greenfield, Bluefield or Brownfield? dealcraftadvisory.ie/insights/erp-journey-greenfield-bluefield-brownfield

DEALCRAFT ADVISORY

Buying ERP, or bidding for it? Make the AI claims comparable.

Dealcraft Advisory provides independent buy-side tender support for organisations procuring ERP, and senior bid direction and independent reviews for the SAP partners bidding to them, drawing on 30+ years of SAP and enterprise deal-making across Europe, the Middle East, Africa and Latin America.

For buyers, we build AI evidence into the requirement and the scoring. For partners, we test your AI story the way an evaluator will, before it reaches one.

Book a free discovery call

For buyers Independent support on ERP procurement, by stage.

Red team reviews Fixed-fee independent review of your proposal, AI claims included.

Bid Director Senior direction of a live pursuit, from qualification to submission.

Try the interactive tools · sales@dealcraftadvisory.ie

Practical guidance based on professional experience, informed by published research. Not affiliated with SAP or any organisation cited. The worked example is illustrative.

Dealcraft Advisory is a registered business name, CRO Business Name No. 793378. Legal details: dealcraftadvisory.ie/legal. © 2026 Dealcraft Advisory.