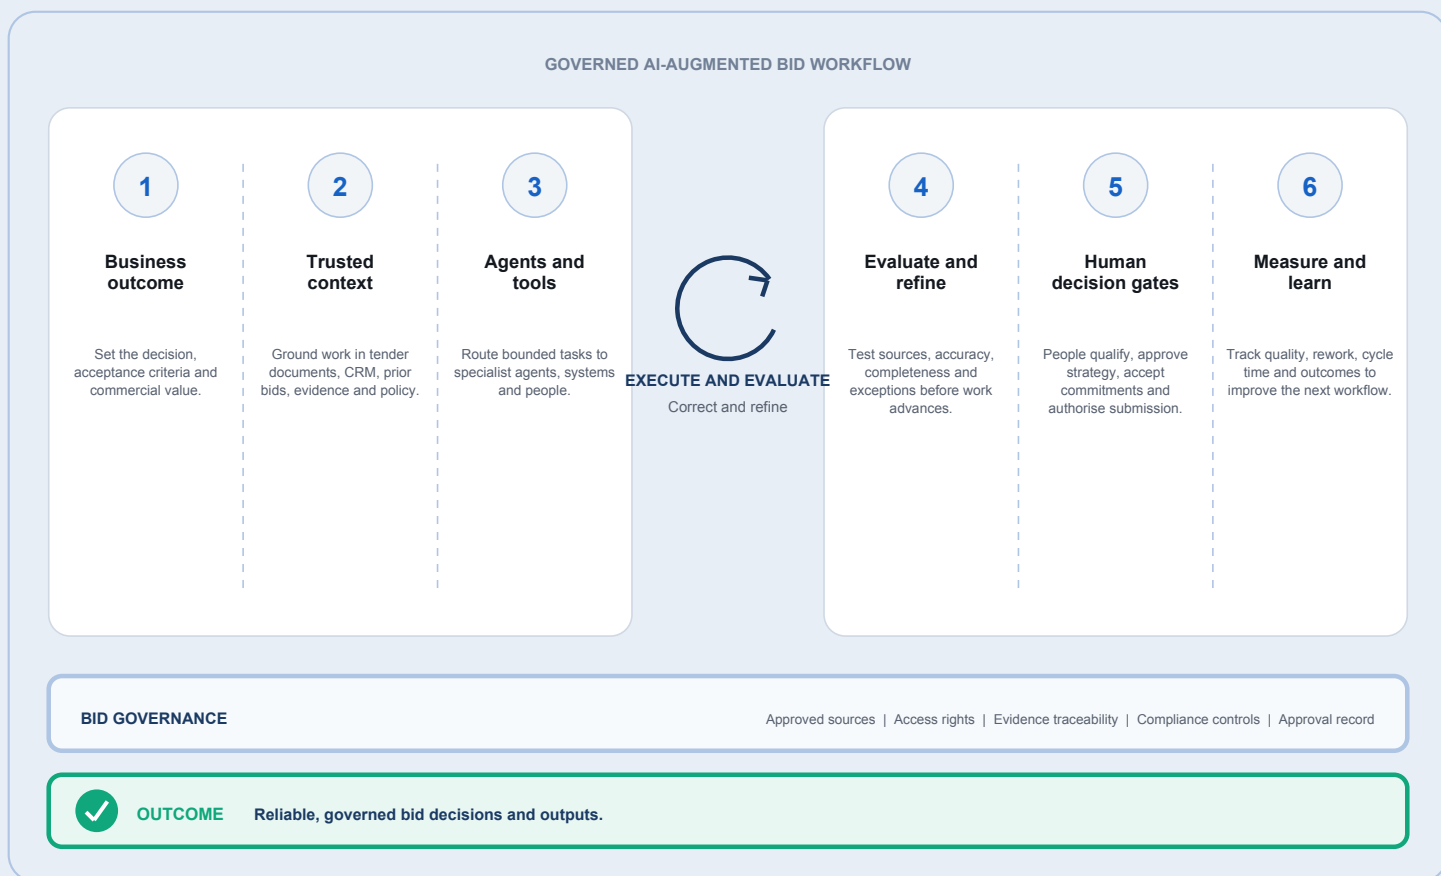


How to build an AI-augmented bid practice

A practical guide to the decisions, context and controls behind reliable human-agent collaboration.



A NOTE FROM YORIK

Build the practice before scaling the tools.

AI can produce more bid work. That does not, by itself, create a stronger bid practice.

Complex bids depend on a chain of commercial, technical and delivery decisions. When AI contributes to that chain, the team needs explicit context, bounded roles, reliable evaluation and named people who remain accountable.

This guide gives you a practical architecture and twelve questions to examine one real workflow. It is not a platform blueprint or a maturity score. It is a way to see what must be true before a controlled AI experiment deserves trust.

HOW TO USE THIS GUIDE

1

Choose one workflow

Use a live, recent or representative pursuit.
Do not rate the organisation in the abstract.

2

Mark what is visible

For each statement, choose Visible, Partial or Unclear and note the evidence behind your answer.

3

Use the gaps

The pattern shows what to resolve before you build, rely on or expand an AI-enabled workflow.

THE CENTRAL REFRAME

Design the decision, context, controls and human gate before choosing the agent.


Yorik Tisseau

Founder, Ignis Leadership | Bid transformation adviser

THE OPERATING MODEL

The architecture of an AI-augmented bid practice

The six elements form one governed loop. Agents prepare and coordinate. People direct, challenge, decide and approve.



1

Start with the outcome

Name the decision, acceptance criteria and commercial value before discussing a tool.

2

Make the boundaries explicit

Define authorised context, the agent's bounded role, evaluation and exception handling.

3

Keep authority human-led

Use decision gates and recorded learning so greater capacity does not weaken accountability.

AREAS 01 AND 02 OF 06

Assess what is visible in one real workflow.

01

DIAGNOSTIC AREA

Business outcome

Keeps AI work tied to a consequential bid or business decision, not a general ambition to automate.

- 1 The workflow begins with a named business or bid decision, not a general ambition to use AI or save time.

MARK ONE

☐

Visible

☐

Partial

☐

Unclear

- 2 The accountable people agree what a useful output must contain and what would make it unsafe or misleading.

MARK ONE

☐

Visible

☐

Partial

☐

Unclear

CHALLENGE QUESTION

Which consequential decision should this work make easier to reach?

02

DIAGNOSTIC AREA

Trusted context

Gives the workflow current, authorised evidence and prevents confident outputs from resting on the wrong material.

- 3 The required tender, customer, commercial, technical and organisational sources are identified, current and available.

MARK ONE

☐

Visible

☐

Partial

☐

Unclear

- 4 The workflow distinguishes approved evidence from stale, unverified, confidential or other-customer material.

MARK ONE

☐

Visible

☐

Partial

☐

Unclear

CHALLENGE QUESTION

Could a reviewer trace every important output to an authorised source?

AREAS 03 AND 04 OF 06

Assess what is visible in one real workflow.

03

DIAGNOSTIC AREA

Agents, tools and roles

Turns a vague request to 'use AI' into a bounded contribution with clear ownership and operating limits.

- 5 The agent contribution is bounded and specific: retrieving, comparing, checking, structuring or routing work.

MARK ONE

☐

Visible

☐

Partial

☐

Unclear

- 6 A named person directs the work, resolves exceptions and decides whether the output is used.

MARK ONE

☐

Visible

☐

Partial

☐

Unclear

CHALLENGE QUESTION

Is the agent's job clear enough without saying 'help with the bid'?

04

DIAGNOSTIC AREA

Evaluation and refinement

Makes failure detectable before it affects the pursuit, and creates a disciplined route to improve the workflow.

- 7 The team can test accuracy, completeness, source use and omissions against a human-reviewed example or baseline.

MARK ONE

☐

Visible

☐

Partial

☐

Unclear

- 8 The workflow defines what happens when evidence is uncertain, sources conflict or the agent fails.

MARK ONE

☐

Visible

☐

Partial

☐

Unclear

CHALLENGE QUESTION

How will you know it is wrong before it affects the pursuit?

AREAS 05 AND 06 OF 06

Assess what is visible in one real workflow.

05

DIAGNOSTIC AREA

Human decision gates

Preserves clear accountability for the choices that shape commercial risk, delivery commitments and submission.

9 Named people retain authority for qualification, strategy, commitments, risk acceptance and submission.

MARK ONE

☐

Visible

☐

Partial

☐

Unclear

10 Each gate states what evidence is required, who approves and how the decision is recorded.

MARK ONE

☐

Visible

☐

Partial

☐

Unclear

CHALLENGE QUESTION

Who can say yes, no or not yet?

06

DIAGNOSTIC AREA

Measurement and learning

Tests whether the workflow improves decisions and capacity, then carries approved learning safely into the next bid.

11 The test measures more than speed: rework, waiting, omissions, corrections and whether decisions happen earlier.

MARK ONE

☐

Visible

☐

Partial

☐

Unclear

12 Approved learning can improve the next pursuit without allowing prior answers or assumptions to become truth.

MARK ONE

☐

Visible

☐

Partial

☐

Unclear

CHALLENGE QUESTION

What could this bid teach the system that the next team can use safely?