MAZi APPLIED INTELLIGENCE

MAZI PROCURA · FOR THE DELOITTE BID LIFECYCLE

Workflow & impact model, in pictures.

Every table in the working draft, redrawn as an infographic — ten stages from bid discovery to win/loss learning, the Win Predictor, and the portfolio view.

Mazi Management Partners LLP

1009, IFCI Apartments, Plot No. 4, Sector 23, Dwarka, New Delhi 110075, India

LLP Identification No. AAH-7683 · mazipartners.ai

Powered by NEORON

Working draft · Confidential · Indicative model

THE JOURNEY AT A GLANCE

Ten stages, one continuous journey

1 — SIGNAL
Discover

Automated bid search — international · national · regional

~100% portal coverage · 5–15 days earlier

2 — DECODE
Qualify

Analyse & categorise to the bid data sheet

Hours, not days · **−80–90%** effort

3 — DECODE + FORGE
Clarify

Pre-bid queries from Deloitte's information strength

−60–70% query-prep time

4 — FORGE
Draft

First-cut technical bid response

−50–70% first-draft effort

5 — FORGE
Differentiate

Best-in-class, doable Approach & Methodology

A&M options **in hours** · marks uplift

6 — SIGNAL + DECODE
Benchmark

A&M and scope vs precedent — first-cut analysis

<24 h vs 3–5 days

7 — RESOURCE FINDER
Resource

Locate & propose suitable experts, with Deloitte precedence

−60–70% longlisting time

8 — WIN PREDICTOR
Predict

Win probability vs Big 4 and non-Big 4 field

Evidence-based **P(win)** before investment

9 — WIN PREDICTOR
Decide

Safe / target / aggressive scenarios on engagement value

QCBS trade-off **in minutes**

10 — PLATFORM
Learn

Win/loss feedback into the models

Compounding accuracy every cycle

Stages 1–7 run the bid · stages 8–10 run the business case around it

STAGE 1 · SIGNAL

Discover — automated bid search

Primary user: **BD analyst / sector leads**

Portals, PSUs, MDB pipelines — crawled, de-duplicated, fit-ranked

	TYPICAL BASELINE	WITH MAZI PROCURA
~100%	**Source coverage** Partial, person-dependent	**Effectively complete** Continuous, all mapped sources
+5–15 days	**Discovery lead time** Often 10–20 days before deadline	**At/near publication** More response runway
−80–90%	**Relevant tenders missed** Material, unmeasured	**Near zero** For mapped sources
−70–80%	**BD monitoring effort** Several hours/day across teams	**Ranked-digest review** No portal-hopping

Indicative modelled ranges — to be validated in pilot

STAGE 2 · DECODE

Qualify — analyse to the bid data sheet

Primary user: **Bid manager**
Every field extracted, source-cited to page & clause

Time to complete data sheet

−80–90%

Baseline — 2–4 person-days

With PROCURA — **< 1 hour + review**

Bid/no-bid decision cycle

−60–70%

Baseline — 3–7 days

With PROCURA — **Same/next day**

Compliance items missed

−50–80%

Baseline — Occasional, costly

With PROCURA — **Checklist-complete, cited**

User journey — upload the bundle → data sheet returns in under an hour → review flagged items only → circulate for bid/no-bid with citations attached.

STAGE 3 · DECODE + FORGE

Clarify — pre-bid queries from information strength

Primary user: **Engagement team**
Each query tied to the clause it targets — and why

−60–70%

Query-pack preparation time

From 1–2 days of senior time to a 2–3 hour review of a drafted, prioritised pack.

1–2 days → 2–3 hours

100%

Risk-clause categories covered

From partial, drafter-dependent coverage to every flagged category addressed.

Partial → all flagged classes

2–3×

Strategic queries per bid

Scope-shaping queries systematically surfaced instead of raised ad hoc.

Few, ad hoc → systematic

User journey — review the drafted, prioritised query pack → edit tone, select which to file → submit before the pre-bid meeting with full clause references.

STAGE 4 · FORGE

Draft — first-cut technical bid response

Primary user: **Bid team**
The single largest cost in the pursuit

Metric	Baseline	With Forge	Change
First-draft assembly effort	5–15 person-days	**1.5–4 person-days**	−50–70%
End-to-end response cycle	3–5 weeks	**2–3 weeks**	−30–40%
Pursuits per quarter, same bench	Baseline n	**n × 1.25–1.4**	+25–40%

Compliance-mapped first draft + gap list — the team's time goes to differentiation, not assembly

STAGE 5 · FORGE

Differentiate — A&M that stands out

Primary user: **Partner / engagement director**
Doable, not decorative — grounded in delivered methods

−70–80%

Time to credible A&M options

2–4 days for one option becomes hours for 2–3 alternative architectures.

1 option in days → 3 in hours

100%

Criteria explicitly addressed

Every evaluation criterion mapped clause-by-clause to text, with an evaluator's view of where each variant earns its marks.

Partial → full mapping

+2–5 /100

Technical-score potential

Criteria-targeted drafting instead of drafter-dependent quality. To validate in pilot.

Marks uplift, targeted

Indicative modelled ranges — to be validated in pilot

STAGE 6 · SIGNAL CORPUS + DECODE

Benchmark — scope vs precedent

Primary user: **Bid manager**

"What did winning look like on tenders like this?"

<24 h

Time to first-cut benchmark

Versus 3–5 days — if done at all. Team shape, effort ranges, methodology patterns.

−80%+ time

5–10×

Comparable tenders considered

Corpus-wide retrieval instead of 2–3 comparables from memory.

Wider precedent base

Every bid

Scope anomalies caught pre-bid

Systematic deviation flags — a pricing and risk signal feeding the query pack and pricing scenarios.

Ad hoc → standard

Anomalies feed Stage 3 queries and Stage 9 pricing

Resource — find & propose suitable experts

Primary user: **Staffing / HR-BD**
Ranked longlist with criteria-fit evidence per candidate

–60–70%

Longlist build time

From 1–2 network-bound weeks to 2–3 days including the start of outreach.

1–2 weeks → 2–3 days

–80%

CV–eligibility fit checking

Automated matrix match against the tender's exact criteria — years, sector, qualifications, comparable projects.

Manual → automated

Every CV

Deloitte-precedence visibility

Prior engagements with the firm systematically surfaced instead of living in tribal knowledge.

Tribal → systematic

STAGE 8 · WIN PREDICTOR — ARBITER LOGIC, REVERSED

Predict — win probability against the field

Primary user: **Partner**
What the partner sees, before major spend is committed

42% ±8

P(win) with confidence band

Under target-case assumptions — before investment.

Score decomposition

Where marks are won and lost:

Eligibility · A&M · CVs · Presentation · Price

Competitor posture map

Likely entrants, their strongest criteria, incumbency and relationship signals.

Sensitivity table

Which single change moves P(win) most — a CV swap, an A&M variant, a price step.

STAGE 8B · SHARED-EXPERT DETECTION

Resource-conflict probability per named CV

- **Historical multi-bid appearances** — high repeat-proposers flagged for exclusivity terms
- **Profile scarcity vs requirements** — scarce profiles = high sharing probability; secure first
- **Engagement-window conflicts** — availability clashes surfaced before committing the CV

Lock exclusivity, substitute, or price the risk into P(win) — before a rival files the same expert.

STAGE 9 · WIN PREDICTOR

Decide — safe / target / aggressive

Primary user: **Pricing partner**

₹-denominated advisory engagement · 80:20 QCBS · illustrative

SAFE CASE · PROTECT MARGIN

30–40%

Modelled P(win)

Tech position: at par with best rival

Price: at internal benchmark

Margin: full margin

TARGET CASE · BALANCED PUSH

45–55%

Modelled P(win)

Tech position: +2–3 marks via A&M + CV upgrades

Price: −4–6% vs benchmark

Margin: modest concession

AGGRESSIVE CASE · MUST-WIN

60–70%

Modelled P(win)

Tech position: +4–5 marks, premium team

Price: −10–12% vs benchmark

Margin: material concession

The QCBS exchange rate — in an 80:20 evaluation, one technical mark conceded costs ≈ a 4% price cut to claw back. The predictor computes this rate for the tender's actual weighting (90:10, 80:20, 70:30, QBS or L1) and shows the price window each technical position buys.

AGGREGATE IMPACT MODEL · PORTFOLIO VIEW

Four headline movements, compounding

−30–50%
Cost per pursuit

Automation of search, data sheet, drafting and resourcing.

+25–40%
Pursuit capacity

Freed senior hours redeployed to more bids — same bench.

+10–20%
Win rate (relative)

Better bid selection via P(win) discipline + criteria-targeted quality.

+40–80%
Revenue per BD rupee

The compound of the three movements — modelled.

The honest caveat

These are modelled ranges, not promises. The credible way to convert them into Deloitte's numbers is a contained pilot — **4–6 weeks, one or two practices, live tenders**, Deloitte's own baseline measured before and after, with success criteria agreed in writing before the pilot starts.