The Municipal BD Intelligence System

A Step-by-Step Replication Playbook

TUJ edition v1.5 — October 24, 2026. *Mirrors the case-study page and the Build Guide; the three are published together.*

How an AI agent system replaced 10–16 hours of weekly manual business-development work with a 15–30 minute review — and how to build the same system in your own market.

Author	Yi Hou, P.E., CFM, PMP — Water Resources practice lead and Project Manager, engineering consulting
Curriculum Adaptor	Temple University, Japan Campus (TUJ) — Applied AI Systems & Enterprise Software Series
Version	Playbook v1.5 — October 2026
Audience	University students replicating the system as an academic case study
Illustrative Market Example	Greater Tokyo Metropolitan Area (Tokyo Waterworks, Tokyo Sewerage, Setagaya City, Yokohama, MLIT)
Companion Document	*AI-Powered Municipal Business Development System — Operational Case Study, Methodology & ROI* (the original deployment)

A note on data and market adaptation. *The production system was originally engineered and proven on live public-sector data. For Temple University Japan (TUJ) students, this playbook illustrates the architecture using the Greater Tokyo Metropolitan Area (Tokyo Bureau of Waterworks, Tokyo Bureau of Sewerage, Setagaya City, Yokohama City, MLIT PPI, etc.) as a concrete classroom case study. What transfers to your classroom is the methodology, prompt design, and system architecture, not any single geography. Follow the same rule in your replication: publish nothing real that you would not put on a public poster.*

For Instructors — Using This Case Study in Class

Learning outcomes. After working through this case study, students can:

- Explain the agent engineering equation — Brain (LLM) + Hands (Tools) + Rules (Instructions) + Memory (State) + Clock (Schedule) — and why a standing instruction pack, not a single prompt, is the real unit of design.
- Distinguish a procedural script, a conversational chatbot, and an autonomous agent, and choose the right pattern for a recurring knowledge-work cycle.

- Design a two-tier evidence pipeline that keeps verified facts separate from early signals, with a named rule (Red Flag) for anything unverified.
- Build a working agent by directing it conversationally: stage the work into steps, write prompts that carry data schemas and guardrails, and verify every step against a written checkpoint.
- Improve instructions empirically — observe a failure mode, write the guardrail that contains it, and fold the rule into the master instructions.
- Assemble staged prompts into one complete set of master instructions, and schedule the system to run its cycle unattended.

Suggested classroom use. Assign this Playbook as pre-reading (the methodology). Run the build hands-on from the companion *Step-by-Step Build Guide* — Steps 1.1–4.1 in four phases, each ending in a checkpoint. Project the live demo dashboard in class: its data is fictional, so students can open and click everything; what they replicate is the deliverable pattern, not the data. The *Technical Addendum* serves advanced students, and the optional rebuild-in-code labs (case-study page §9) extend the case for programming courses.

Assessment. Every step's checkpoint is written as a pass/fail test. The checkpoint artifacts — the positioning statement, the portal registry, the system file, and the final master prompt — can be collected as staged submissions, so grading follows the build itself.

Citation. Hou, Yi. *Building an Autonomous Municipal BD AI Agent.* TUJ Case Study Series, Temple University, Japan Campus, 2026. Web edition v1.5 — October 24, 2026. https://yihou.dev/case-study/

Reuse. Instructors may assign, link to, and print this case study and its PDFs for classroom use, with attribution. Please link to the living page rather than redistributing file copies, so students always land on the current edition.

Contact. Questions from the class are welcome, and the builder is happy to join a Q&A session by video call. Contact: hymonkeyking@gmail.com

1. The Tool at a Glance

Definition. The Municipal BD Intelligence System is an autonomous AI agent system that runs the weekly business-development (BD) intelligence cycle for a municipal engineering consulting practice: it monitors public procurement portals, tracks opportunities *before* they are advertised, maintains one-page account strategies ("Gold Sheets") for priority clients, profiles key decision-makers, filters networking events, and delivers everything as a single interactive dashboard on a fixed weekly schedule.

Who built and operates it. The system was designed and built by practicing engineers — a water-resources lead and a data-technology lead — not by a marketing department and not by software developers. It was built by *directing an AI agent* through structured instructions, templates, and rules; no hand-written application code was required. This is the central

lesson of the case study: **a domain expert who can articulate a methodology can build the tool that executes it.**

Build stack (as built).

Layer	What was used
Agent	A Gemini Spark agent (a current Gemini model — names change; see the case-study page, Step 1.1, Platform facts)
Interface	A standalone, single-file HTML dashboard (100% English), generated by the agent
Persistence	Browser-local storage in the dashboard, with one-click JSON backup/export
Cadence	Full pipeline re-run once a week, on a fixed schedule
Archive	Every cycle's dashboard archived to Google Drive (`Gemini Agent > Marketing > Output`)
Push	The weekly dashboard is emailed to the operator automatically
Methodology	Miller Heiman Strategic Selling® — account-level Gold Sheets (LAMP), Buying Influences, Red Flags, Win-Results

Market context. Municipal water/civil engineering consulting in large urban clusters: dozens of public bodies buy engineering services through public proposals and tenders, and the winners are usually decided by relationships and positioning established *months before* a solicitation is posted.

2. The Goal

Primary goal. *Never be surprised by a municipal opportunity again.* See every relevant solicitation early, understand every priority account on one page, and walk into every meeting with a prepared strategy — while cutting the weekly BD effort from 10–16 hours to under 30 minutes.

The goal decomposes into four sub-goals, one per structural bottleneck of traditional municipal BD:

#	Bottleneck in the traditional workflow	Design goal of the system
G1	**Portal fragmentation & scraping blind spots.** Each municipality uses a different procurement platform (Tokyo e-Procurement, e-Tokyo Joint, MLIT PPI in the Tokyo classroom	Monitor every target portal with a browser-capable agent, on a 45–60 day lookahead horizon, so no in-scope solicitation is ever missed.

	example); several are JavaScript-rendered single-page apps that generic crawlers cannot read. Engineers burned 4–6 hours/week checking portals manually — and still missed major solicitations.	
G2	**The pre-RFP disconnect.** Once an RFQ or tender is posted, incumbents may have spent 6–12 months shaping the scope, and a statutory "cone of silence" (or nyusatsu kenshou) restricts contact. Waiting for the posting is too late.	Detect opportunities 30–90 days *before* advertisement from planning signals (capital improvement plans, 5-year facility renewal plans, contract expirations, appointments) and run disciplined pre-positioning.
G3	**Decision-maker information gap.** Teams approached agencies with generic capability statements, with no map of who actually decides (Economic, User, and Technical Buyers) and where the firm's credentials align.	Maintain a verified, six-dimensional profile of every key decision-maker, including credential overlaps (dual PE/CFM or Gijutsushi) and a ready-to-use ice-breaker.
G4	**Networking friction.** Dozens of association events (JSCE, JWWA, MLIT-hosted forums in Japan) announced through scattered mailing lists; engineers either traveled badly or missed the rooms that matter.	Filter events to a 45-minute travel radius and a rolling 14-day window, and generate a one-page "battlecard" before every pivotal meeting.

3. Background: The Methodology — Strategic Selling and the Gold Sheet

The tool in this playbook is not a generic "AI dashboard." Its operating logic comes from a specific, published sales methodology with a forty-year history. Students who understand where the methodology comes from will build better replications than students who only copy the dashboard's layout.

3.1 Origins: Miller, Heiman, and the complex sale

In the late 1970s, American sales consultants **Robert B. Miller** and **Stephen E. Heiman** were studying a recurring failure pattern in large business-to-business sales: deals involving many decision-makers were being lost not on product merit, but because sellers negotiated with the wrong people, never identified their internal opponents, and never learned what each

decision-maker would personally count as a win. Their answer was a planning discipline for the **complex sale** — a sale in which no single person can say yes alone.

They published the method in 1985 as ***Strategic Selling**, and revised it — with Tad Tuleja — as* **The New Strategic Selling** *(1998; updated edition 2005). The books became standard texts in enterprise sales training, taught for decades through the Miller Heiman organization — today part of Korn Ferry**, which continues to deliver the programs.

3.2 The core concepts — the vocabulary the tool runs on

- **Buying Influences.** Four roles exist in every complex decision: the **Economic Buyer** (the only role that can approve the money), the **User Buyer** (will live with the result day to day), the **Technical Buyer** (screens against requirements and can veto, but cannot approve), and the **Coach** (an insider who wants you to win and will teach you how). These are *roles, not titles*: one person can hold several, and one role can be shared by several people.
- **Response Modes.** How a buyer currently feels about change: **Growth** (funded and ambitious — talk vision), **Trouble** (something hurts — talk remedy), **Even Keel** (content — the hardest to move), **Overconfident** (believes there is no problem). The same message lands or fails depending on the mode it meets.
- **Red Flags.** Any missing information or adverse signal in your position. "We have never met the Economic Buyer" is itself a Red Flag — the discipline is to surface unknowns, then make eliminating the largest one your first action.
- **Win-Results.** Clients do not buy features; each Buying Influence buys a *result* that is a win for them, personally and professionally. A pitch is complete only when it can be stated as a Win-Result for every influence.
- **Ideal Customer Profile & the funnel.** Decide in advance which accounts deserve pursuit at all, and grade positions honestly — *above the funnel* (known, unengaged), *in the funnel* (active dialogue), *best few* (where deep resources go).
- **The Win-Win rule.** A sale that is not a win for the buyer is treated as a defective outcome, not a victory: it will not survive implementation or renewal.

3.3 Three methods, three sheets — and a printer story

The Miller Heiman firm grew the original method into a family of three, each planning at a different level, each with a one-page worksheet whose paper color became its name:

Method	What it plans	Horizon	Its worksheet
Strategic Selling	A single opportunity	One deal	**Blue Sheet**
Conceptual Selling	A single customer meeting	One meeting	**Green Sheet**
LAMP — Large Account Management Process	A strategic account	1–3 years	**Gold Sheet**

By the well-known account, the colors are an accident of print-shop history: in 1978, a printer in Reno, Nevada, producing the firm's workshop worksheets ran out of white paper and

printed them on blue stock. "Blue Sheet" stuck — and the Blue Sheet remains the family's most famous artifact, still used in executive deal reviews today. The later sheets were named in the same spirit.

3.4 The Gold Sheet, precisely

The Gold Sheet belongs to **LAMP**, the account-level method. Where Strategic Selling asks *"how do we win this deal?"*, LAMP asks *"where is this account going over the next one to three years, what is our position in it, and what must we do — and stop doing — to become its strategic partner?"* Its working parts, as taught in the LAMP program, include: a deep account analysis (key players; trends and opportunities in the account's "field of play"; your organization's strategic strengths; and your **critical vulnerability**), a **charter statement** for the account, relationship goals, **focus investments** and **stop investments**, revenue targets, and the single sales objectives that follow from them.

Two properties matter for this case study. First, a Gold Sheet is a **living document** — reviewed and revised as account evidence changes, the opposite of an annual plan written once and filed. Second, it is deliberately **limited to the largest, most strategic accounts**: the analysis is expensive, and spreading it over every account destroys it.

A precision worth teaching. The words "Gold Sheet" do not appear in *The New Strategic Selling* at all — verifiable by a full-text search of the book. The book teaches Strategic Selling and its Blue Sheet; the Gold Sheet is LAMP's tool, delivered through the LAMP program. Secondary writing frequently blurs the two. This playbook keeps them separate, and deliberately combines them: the system in §1 runs account-level **Gold Sheets** (LAMP's level) using Strategic Selling's **concepts** (Buying Influences, Response Modes, Red Flags, Win-Results) as their working vocabulary.

3.5 Why a methodology from 1985 runs an AI system from 2026

Three properties made this framework the right operating system for an autonomous agent:

1. **It is explicit.** Every judgment the method requires — who decides, what mode they are in, what is missing — is a named field. Named fields can become a data model, a template, and a checklist an agent must complete.

2. **It is honest by construction.** Red Flags force unknowns onto the page. That is exactly the discipline that keeps a generative agent from confabulating a pipeline — uncertainty is recorded, never papered over (see guardrail R1 in §7).

3. **It is level-aware.** Public-sector engineering pursuits are won at the *account* level over years, not at the *opportunity* level over weeks. Choosing the Gold Sheet rather than the Blue Sheet as the system's central artifact was a deliberate design decision (§4, D2/D4) — and students replicating this system in another industry should make the same level choice consciously, before writing a single prompt.

3.6 Background reading

- Miller, R. B., & Heiman, S. E. (1985). *Strategic Selling*. New York: Warner Books. — The original statement of the method.
- Miller, R. B., Heiman, S. E., & Tuleja, T. (1998). *The New Strategic Selling* (rev. ed.; updated edition 2005). New York: Warner Books. — The edition studied for the original deployment; the source of the concept summaries in §3.2.
- Miller, R. B., Heiman, S. E., & Tuleja, T. (2003). *The New Conceptual Selling*. New York: Warner Business Books. — The meeting-level companion method (the Green Sheet).
- Korn Ferry. *Large Account Management Process (LAMP)* — program overview: https://www.kornferry.com/capabilities/leadership-professional-development/training-certifications/large-account-management-process — The current official description of LAMP and its Gold Sheet analysis, from the firm that now owns the Miller Heiman methodologies.
- Korn Ferry. *Strategic Selling® with Perspective* — program overview. — The current training form of Strategic Selling, for students who want the methodology's modern presentation.

4. How It Was Conceived — Seven Design Decisions

The system was not designed "AI-first." It was designed **workflow-first**, and each design decision answers a specific failure of the manual process.

- **D1. Start from the calendar, not the technology.** The product is a weekly cycle — scan, verify, update strategy, prepare meetings, log actions. The dashboard is simply the visible output of that cycle. If you cannot state your weekly ritual in one sentence, you are not ready to automate it.
- **D2. Methodology before automation.** The agent does not "search the web for leads." It operates a named, published sales methodology — Miller Heiman Strategic Selling: every priority account gets a Gold Sheet; every account analysis names the four Buying Influences (Economic Buyer, User Buyer, Technical Buyer, Coach); every gap is logged as a Red Flag; every pitch is expressed as a Win-Result (what the client gains), never as a feature list. The methodology is the agent's operating system. (Background: §3.)
- **D3. Two-tier evidence discipline.** All opportunities live in exactly one of two tiers, and the tiers never mix:
- **Tier 1 — Verified Live Solicitations:** active postings only, each with a confirmed solicitation ID, a verified deadline, and a direct link to the official posting.
- **Tier 2 — Pre-RFP Strategic Pursuits:** early intelligence (a CIP allocation, a facility renewal plan, an expiring master contract, a new appointee) with a catalyst and a timing window — and *never a fabricated deadline*.
- **D4. Pre-positioning beats bid-chasing.** Because of the cone of silence (and anti-collusion rules), the system's center of gravity is Tier 2. A solicitation appearing in Tier 1 is the *end* of a pursuit that should have started months earlier, not the beginning of one.

- **D5. Reposition around a differentiator, not a brochure.** The practice repositioned itself as **"Digital Water & Asset Management"** — a dual-engine offer pairing a data-technology group (AI analytics, municipal web portals, utility data systems) with a water-resources group (stormwater, hydraulics & hydrology, floodplain). Every Gold Sheet's strategy section is written against this positioning, so the system sells an integrated outcome instead of generic civil design hours.
- **D6. The agent prepares; the human decides and acts.** The system never contacts anyone. Go/no-go decisions, meeting requests, and every external message remain human acts. This is both an ethics rule and a quality rule: the operator's judgment is the final verification gate.
- **D7. Ship a single-file artifact.** The dashboard is one standalone HTML file with local persistence — no server, no IT ticket, no login wall between the operator and the information. A fictional-data twin of the dashboard serves as the shareable demo, so the methodology can be shown publicly without exposing real business data.

5. Inputs

The system consumes **public information only**. Every input below is a published source; the firewalls in §7 keep everything else out.

#	Input	Typical sources (classroom example)	Cadence	Notes
I1	Live solicitation postings	Municipal procurement portals (Tokyo e-Procurement, MLIT PPI in Japan)	Weekly scan	Many portals are JavaScript-rendered; a browser-capable agent is required — plain web crawlers miss them (bottleneck G1).
I2	Capital improvement plans & budgets	City CIP documents, 5-year facility renewal budgets (Shisetsu Seibi Keikaku), bond filings	Weekly scan for changes	The primary Tier 2 signal: money allocated today becomes an RFQ/tender in 1–3 quarters.
I3	Council agendas & committee minutes	City council and utility board published agendas and assembly minutes	Weekly	Where initiatives, contract renewals, and study authorizations first surface.
I4	Incumbent contract data	Prior award records,	On change	An expiring incumbent

		master/on-call consultant registries		contract is a dated, predictable pursuit window.
I5	Personnel & appointment announcements	City news releases, official gazettes, association publications	Weekly	A new appointee's first 90 days are the highest-value outreach window in this market.
I6	Public professional profiles	Published bios, conference speaker pages, association rosters	At profiling time	Public sources only; feeds the decision-maker dossiers in §6.
I7	Event calendars	JSCE, JWWA, JSWA, MLIT in Japan	Weekly	Filtered by the radius/window rules in §7 (R7).
I8	Internal capability inventory	The firm's own service descriptions, group leads, project history	Maintained by operator	Used to match each opportunity to a lead capability and write Win-Results honestly.

6. Outputs

One weekly cycle produces one dashboard plus derived working documents. The dashboard's modules, top to bottom:

#	Output	What it contains	Who consumes it
O1	KPI header	Counts at a glance: portals monitored, active verified solicitations, pre-RFP strategic targets, Gold Sheets maintained, and a "golden outreach window" indicator naming the account whose timing window is open now.	Operator, practice leadership
O2	Tier 1 table — Active Bids & RFQs	One row per verified solicitation: Agency · Solicitation ID & Scope · Verified Deadline · Procurement Portal · Direct Posting Link.	Operator (pursuit decisions)
O3	Tier 2 table — Pre-RFP Target Pipeline	One row per early pursuit: Target Agency ·	Operator (positioning

		Strategic Project Initiative · Pre-positioning Catalyst & Timing · Lead Capability · Procurement Portal Base · Strategic Action · Source (link + date).	plan)
O4	Strategic Gold Sheets	A one-page account playbook for each priority account (six in production), printable, updated cumulatively — a long-term asset, not a weekly snapshot. Sections in Appendix B.	Operator, pursuit team
O5	Decision-maker dossiers	One profile per key official, in six dimensions (§8, Step 2.2), fact-checked against public sources and linked to the account's Gold Sheet.	Operator, before any contact
O6	Event radar	Networking events inside the radius/window rules, each tagged with the audience it draws and a direct registration link.	Operator
O7	Pre-event battlecards	A one-page brief before pivotal meetings: who will be there, their technical pain points, and the strategic questions to ask.	Operator (mobile, at the meeting)
O8	Weekly execution log	Three editable blocks — *What was accomplished last week · Action items for next week · Strategic observations & feedback notes* — saved in browser localStorage.	Operator, practice leadership
O9	Weekly push & archive	The regenerated dashboard emailed to the operator after each weekly run and archived to Google Drive, creating a permanent,	Operator, leadership

		dated history of every cycle.	

7. Restrictions & Guardrails

These rules are part of the design, not afterthoughts. Most of the system's credibility — the reason leadership trusted its output — comes from R1–R3.

- **R1. Verified-or-labeled, always.** A Tier 1 item must have a confirmed solicitation ID, a deadline verified against the official posting, and a direct link. A Tier 2 item must be labeled as early intelligence and must never carry an invented date. If a fact cannot be verified, the dashboard shows the gap explicitly (in Gold Sheet language: it becomes a **Red Flag**) rather than filling it in.
- **R2. Public information only.** The system ingests published sources only. Personal, non-work contacts are firewalled out of the system entirely.
- **R3. Read-only on people.** Profiles are assembled exclusively from public material. The system never sends a message, a connection request, or any other contact. All outreach is drafted for the operator, who alone decides whether and how to send it.
- **R4. Procurement-ethics compliance.** Municipal solicitations operate under a cone of silence (or nyusatsu kenshou) once posted: no positioning contact with agency staff about a live procurement. The system's response is structural — do the positioning work in Tier 2, *before* posting, and use only official channels afterward.
- **R5. No confidential data in shared artifacts.** Anything shown outside the practice (demos, teaching materials, this playbook's classroom examples) uses fictional or public benchmark data and is labeled as such.
- **R6. Gold Sheets only for the Best Few.** A full Gold Sheet is expensive thinking. The system maintains them only for the priority accounts (six in production). Every other account gets a one-line memo. Depth where it matters beats shallow coverage everywhere.
- **R7. Event discipline.** Events enter the radar only if they fall within a 45-minute travel radius of the office/campus and inside the rolling 14-day window. (The radius deliberately includes regional coordination hubs, e.g., regional government centers in the Kanto region.)
- **R8. Accepted MVP constraints.** Persistence is browser-local storage plus manual JSON backups; the system is single-operator by design, with no multi-user backend and no role-based access. These are known, accepted limitations of the first version — documented here so students improving the design start from honesty about them.

8. Step-by-Step Replication Guide

Each step lists its objective, the actions to take, and a **checkpoint** — the test that tells you the step is done. Steps 1.1–2.2 build the system's content; Steps 2.3–2.4 build and automate the container; Steps 2.5–4.1 make it trustworthy and adopted.

Step 1.1 — Define the operator and the weekly ritual

- **Objective:** Fix the two things every later decision depends on: who runs the system, and what the weekly cycle is.
- **Actions:** Name one operator (in the original: the practice lead who owns BD for the office). Write the weekly ritual in one sentence — e.g., "Every week on run day, scan all monitored portals, verify new items, update the six Gold Sheets, prepare next week's meetings, and log the week's actions." Choose a fixed run day/time; the discipline of a fixed cadence matters more than the specific slot.
- **Checkpoint:** You can state, in one sentence each: the operator, the run day, and the five verbs of the cycle (scan → verify → strategize → prepare → log).

Step 1.2 — Build the portal registry

- **Objective:** A complete, explicit list of where opportunities in your market are published. This registry is the system's coverage boundary — anything not in it does not exist for the system.
- **Actions:** Enumerate your target agencies. (Classroom Tokyo example: Tokyo Bureau of Waterworks, Tokyo Bureau of Sewerage, Setagaya City, Shinjuku City, Yokohama City, MLIT Kanto Bureau.) For each, record: agency name, procurement portal platform and URL (e.g., Tokyo e-Procurement System, e-Tokyo Joint Service, MLIT PPI), and categories in scope. Flag which portals are JavaScript-rendered single-page apps — these dictate tooling: you will need an agent that operates a real browser, because conventional crawlers return empty pages from them.
- **Checkpoint:** Every in-scope agency has a registry row with a working portal link you have opened yourself at least once.

Step 1.3 — Define the two-tier data model

- **Objective:** Encode the evidence discipline (D3/R1) as data structure, so the tool cannot silently mix a rumor with a verified posting.
- **Actions:** Create the two tables with exactly these fields:
- **Tier 1 — Solicitation:** Agency · Solicitation ID · Scope (one line) · Verified Deadline (date + time) · Portal · Direct Posting Link · Fit note.
- **Tier 2 — Pre-RFP Target:** Target Agency · Strategic Initiative · Catalyst & Timing (the event that opens the window — a CIP allocation, a facility renewal plan, a contract expiry, an appointment — with its date or quarter) · Lead Capability (internal group) · Portal Base · Strategic Action · Source (URL, date, and the original quoted line that supports the timing).

- Rule: a Tier 2 row may not contain a deadline field *at all*. When a Tier 2 pursuit is advertised, it *graduates*: verify it and re-enter it as Tier 1.
- **Checkpoint:** Show your schema to a skeptic and have them try to insert a rumored deadline. If the structure accepts it, the schema is wrong.

Step 1.4 — Select the Best Few accounts and write the first Gold Sheets

- **Objective:** Choose the small set of accounts that justify deep, cumulative strategy work, and encode the methodology (D2) as a reusable one-page template.
- **Actions:** Score candidate accounts on: sustained capital budget, problems matching your differentiator, an identifiable decision chain, and a timing catalyst (expiring incumbent contract, new leadership, new funding). The original selected six priority accounts (illustrated in the classroom case as: TMG Waterworks, TMG Sewerage, Setagaya City, Yokohama Waterworks, Hachioji City, and Kawasaki City). For each, complete a Gold Sheet (template in Appendix B): account overview and response mode (Growth / Trouble / Even Keel / Overconfident); the four Buying Influences with a Win-Result for each; competitive landscape and vulnerabilities; sales strategy; and an action plan whose first item is always *eliminate the largest Red Flag*.
- **Checkpoint:** Each Gold Sheet fits on one page, names at least one human in at least two of the four Buying Influence roles (with unknowns flagged), and ends with a dated action list.

Step 2.1 — Encode the fit and prioritization rules

- **Objective:** Decide in advance what deserves attention, so weekly scanning is classification, not deliberation.
- **Actions:** Write explicit rules: which service lines are core (floodplain modeling, stormwater master planning, hydrology & hydraulics, utility rehabilitation, and applied data-technology tools), which are adjacent (review case-by-case), and which are out of scope (ignore silently). Add the lookahead rule: Tier 1 items are worked on a 45–60 day horizon — anything closing sooner gets an escalation flag, not a shrug.
- **Checkpoint:** Hand the rules to another person; they should classify five sample postings and agree with your classification on at least four.

Step 2.2 — Build the decision-maker dossier template

- **Objective:** Turn "know your client" into a repeatable six-dimensional profile that can be produced in minutes from public sources.
- **Actions:** For the lead official on each active pursuit, complete:

  1. Responsibility & Role — what they own, who they report to, which Buying Influence they are.

  2. Background & Pedigree — career path, credentials, prior employers (from published bios).

3. Core Similarities (Leverage) — genuine overlaps with your team (e.g., dual PE/CFM or Gijutsushi credentials, university alumni networks — record only overlaps that are *true*).

4. Conferences & Gatherings — where they appear publicly; the rooms where a legitimate introduction can happen.

5. High-Frequency Keywords — the terms and priorities in their public statements; the language of their agenda.

6. Ready Ice-Breaker Pitch — a three-sentence, technically specific opener built from dimensions 1–5, ready before any meeting.

- Mark every claim with its source; mark anything unverified as such. (Guardrail R3 applies: this is reading, never contact.)
- **Checkpoint:** Pick any dossier at random: every factual claim traces to a public source, and the ice-breaker could be said aloud without embarrassing anyone.

Step 2.3 — Assemble the dashboard (build with an agent, not by hand)

- **Objective:** Produce the single-file dashboard that renders Steps 1.2–2.2 as one screen of truth.
- **Actions:** Direct your AI agent module by module, in the output order of §6: KPI header → strategy/positioning banner → Tier 1 table → Tier 2 table → Gold Sheets (with a selector and print view) → dossiers → event radar → weekly log. Give the agent: the data schema (Steps 1.3/1.4/2.2), real sample content, and the guardrails (notably: Tier 2 has no deadline column). Require: editable weekly-log blocks, local persistence, a one-click JSON backup, and a full English interface. Iterate visually.
- **Checkpoint:** Open the file on a machine that has never seen your project. Everything renders from the single file; the log edits save and survive a page reload; the JSON backup restores them.

Step 2.4 — Automate the weekly cycle

- **Objective:** Convert the manual ritual into a scheduled agent run with a human verification gate.
- **Actions:** Schedule the agent (e.g., one fixed slot each week) to: re-scan the portal registry (I1–I7), draft updates to both tiers, refresh dossiers for accounts with open windows, rebuild the dashboard, archive the dated copy to the shared drive, and email it to the operator. The operator then spends 15–30 minutes reviewing: accept, correct, or reject each change (the verification gate, Step 2.5) and finalize the week's action list in the log.
- **Checkpoint:** Run one full cycle end-to-end on schedule, with the operator only reviewing. Time the review. If it exceeds ~30 minutes, the agent's drafts — not the operator — need more work.

Step 2.5 — Install the verification gate (never skip this)

- **Objective:** Make "verified" a claim the system has earned, week after week.

- **Actions:** Before any item may appear as Tier 1, a human (or the agent acting under human spot-check) confirms ID, deadline, and link against the official posting itself — never against an aggregator's copy. Institute a monthly audit: sample ten dashboard items at random and re-verify them at the source; log the error rate. The system treats a single fabricated or stale item as a process failure to be root-caused, not a typo to be fixed quietly.
- **Checkpoint:** One completed audit with a written error count of zero — or a written corrective action.

Step 3.1 — Build the fictional-data demo twin

- **Objective:** Make the methodology shareable — for leadership demos, teaching, and recruiting — without exposing live business data (R5).
- **Actions:** Duplicate the dashboard; replace every agency, person, solicitation, and figure with invented ones (or classroom case study examples); label it "DEMO — Fictional Data" on every view. Verify by search: no confidential production name or number survives anywhere in the file.
- **Checkpoint:** A colleague outside the practice can use the demo to explain the whole methodology — and could not learn a single real fact about your live proprietary pipeline from it.

Step 4.1 — Roll out and iterate

- **Objective:** Convert a personal tool into a practice asset.
- **Actions:** Demo with the twin; collect leadership feedback; expand the portal registry and account list only as fast as the verification gate stays honest; feed every lost pursuit and every stakeholder correction back into the fit rules (Step 2.1) and Gold Sheets (Step 1.4). In the original deployment, this is the step where the system traveled upward: demonstrated to senior executives and program managers, and adopted as the office's standard BD review artifact.
- **Checkpoint:** Someone other than the builder uses the dashboard to run one weekly review unaided.

9. The Intended End State

When the system is working, the operator's week looks like this:

- **On run day** — the new dashboard arrives by email; the dated copy is already archived.
- **A 15–30 minute review** — check the KPI deltas (what is new, what is closing, whose window opened), decide *pursue / watch / pass* on each new item, confirm the Gold Sheet action list, and glance at next week's battlecards.
- **During the week** — meetings happen with dossiers and battlecards on the operator's phone; outcomes are typed into the execution log as they happen; the next cycle's agent run picks them up.

- **Continuously** — every priority account's strategy lives on one page, current to the week; every deadline in the system is one click from its official source; no opportunity in scope is discovered by accident.

The deeper end state is positional, not clerical: the practice walks into procurements having *already been useful* to the client for months — the "trusted advisor before the posting" position that the cone of silence makes impossible to buy at the last minute.

10. How Success Is Measured

Success is measured on two levels: **efficiency** (did the system give the hours back?) and **integrity/coverage** (can its output be trusted, and does it miss anything?). A system that is fast but wrong has failed; a system that is right but unused has also failed.

10.1 Efficiency — Results Reported from Production Deployment

BD workflow stage	Traditional manual process	With the system	Gain
Multi-portal bid scanning	4–6 hours / week	Automated	100% automated; 45–60 day horizon
Decision-maker profiling	2–3 hours / pursuit	< 5 minutes	≈95% saved
Gold Sheet upkeep	2–3 hours / account	Interactive, saved in-browser	≈90% saved
Event & route filtering	1–2 hours / week	Automated	100% automated
Pre-meeting battlecard prep	1–2 hours / meeting	Generated instantly	≈90% saved
Total weekly BD time	**10–16 hours / week**	**15–30 minutes / week**	**>95% freed; 500+ senior engineering hours reclaimed/yr**

10.2 Coverage & Integrity Metrics (Track These in Your Replication)

Metric	Definition	Target
Tier 1 capture rate	Verified in-scope solicitations appearing in the dashboard ÷ in-scope solicitations actually published on monitored portals (monthly audit)	100%
Detection latency	Days between a posting's publication and its appearance in the dashboard	≤ 7 (one cycle)
Lookahead horizon	Days between detection and deadline for Tier 1 items	≥ 45
Pre-positioning lead time	Days between first Tier 2 engagement and the eventual	30–90

	posting	
Gold Sheet coverage	Priority accounts with a Gold Sheet updated within the current cycle	6 / 6 (100%)
Dossier completeness	Active-pursuit dossiers with all six dimensions filled from sourced facts	100%
Data integrity	Fabricated, mislabeled, or stale items found in the monthly audit	0
Review time	Operator minutes to complete the weekly verification & decision review	≤ 30

10.3 Adoption — The Qualitative Test

The original deployment's decisive success signal was organizational, not numerical: stakeholders above and beside the builder — the data-technology SVP, the program manager, and the regional marketing director — asked to use the system and its outputs in their own planning. Measure adoption explicitly: **number of stakeholders who run a review from your dashboard without you in the room.** Zero is a failing grade no efficiency number can offset.

11. Limitations & Failure Modes (Read Before You Build)

- **Portal fragility.** Portals change platforms, add login walls, and restructure pages without notice; a scanner that worked in January silently returns nothing in March. Mitigation: the KPI header makes silence visible (a portal count that stops changing is a warning), and the monthly audit samples the sources, not just the dashboard.
- **Aggregator rot.** Third-party bid aggregators lag and garble deadlines. They are leads, never evidence — hence R1's verify-at-the-official-posting rule.
- **Hallucination pressure.** A generative agent asked for a deadline it cannot find is tempted to supply a plausible one. The two-tier schema (Step 1.3), the no-deadline-field rule for Tier 2, and the verification gate (Step 2.5) exist precisely because this failure mode is *certain*, not hypothetical.
- **Methodology decay.** Gold Sheets rot if they are not updated after every meeting and every new fact; a stale Gold Sheet is worse than none, because it is believed. The weekly log is the forcing function — if it isn't written, the sheets aren't either.
- **Single-operator dependency.** The MVP is one person's system: local persistence, personal scheduling, personal judgment at the gate. Documenting the method (this playbook) is the mitigation; a multi-user backend is the obvious v2.
- **Local-storage risk.** Browser-local data can be lost with a cache clear. Mitigation: the JSON backup ritual at the end of every cycle, plus the dated Drive archive of every dashboard.

- **Scope creep.** Every team wants its accounts added. R6 (Best Few only) is what keeps the deep work deep. Resist.

Appendix A — Data Model (Field Reference)

Solicitation (Tier 1): Agency · Solicitation ID · Scope (one line) · Verified Deadline · Portal · Direct Link · Fit (Core / Adjacent / Out) · Owner · Status.

Pre-RFP Target (Tier 2): Target Agency · Strategic Initiative · Catalyst & Timing · Lead Capability · Portal Base · Strategic Action · Source (link, date, quoted line) · Owner · Window status. *(No deadline field — by design.)*

Gold Sheet (per account): Account Overview · Response Mode · Four Buying Influences (name, role, stance, Win-Result each) · Competitive Landscape & Vulnerabilities · Sales Strategy & Win-Results · Action Plan & Red Flags · Last-updated date.

Decision-Maker Dossier: Name · Title · Agency · Buying Influence role · the six dimensions of Step 2.2 · Sources · Last-verified date.

Event: Name · Date/Time · Location & travel time · Organizer · Target audience · Link · Linked account (if any).

Log Entry (weekly): Accomplished (list) · Next actions (list, with owners) · Strategic observations & feedback notes.

Appendix B — The Gold Sheet on One Page (Template)

1. Account Overview & Response Mode — the agency in one paragraph: size, budget/CIP scale, current priorities; and its mode — *Growth* (funded and ambitious: talk vision), *Trouble* (a failure, a fine, a flood: talk remedy), *Even Keel* (stable: talk optimization and risk), *Overconfident* (believes it has no problem: build trust in small steps first).

2. Four Buying Influences — Economic Buyer (approves the money), User Buyer (will live with the result), Technical Buyer (can veto on compliance), Coach (an insider who will teach you how to win — earned, never assigned). For each: name, stance toward you, and their Win-Result.

3. Competitive Landscape & Vulnerabilities — incumbents and rivals, their exposed flanks (expiring contracts, thin local benches), and your genuine moat. Public, verifiable facts only; never disparagement.

4. Sales Strategy & Win-Results — the one-sentence strategy and the three messages, all grown from the client's agenda (their CIP, their pain, their new leader's first-100-days plan), never from your brochure.

5. Action Plan & Red Flag Elimination — dated next actions with owners; the first action always attacks the largest Red Flag.

Appendix C — Classroom Exercises

1. Rebuild the registry (Week 1). Choose an urban cluster and a professional-services sector (e.g., Tokyo Metropolitan waterworks consulting; or public-sector IT services for a prefectural government). Produce the Step 1.2 portal registry and a populated Tier 1/Tier 2 table set for one weekly cycle. Deliverable: the two tables plus a one-paragraph coverage statement — what your system *cannot* see.

2. Write a real Gold Sheet (Week 2). Pick one real public agency in your chosen market. Using public sources only, complete the Appendix B template for your hypothetical firm. Every unverifiable item must appear as a Red Flag with a proposed elimination action. Deliverable: one page, plus a source list.

3. Break the system (Week 3). You will be given a sample dashboard containing three planted defects: a wrong deadline, a fabricated Tier 2 date, and an unsourced dossier claim. Find all three, name the guardrail each one violates (R1–R8), and write the process change that would have caught it. Deliverable: a one-page audit memo.