

THE UNMANAGED ADVANTAGE OF AI

*Why AI Governance — Not AI Tools — Is the Decision in Front
of Every Economic Development Organization*

WHITE PAPER | JULY 19, 2026

Devin Hillsdon-Smith

Founder & Principal | Hyphen Strategies, LLC

This white paper presents a governance-first framework for artificial intelligence in economic development organizations — how to identify the exposure your organization already carries, how to interrogate a vendor claim before you believe it, how to tell which workflows actually return staff time, and how EDOs can become machine-readable to the systems now shaping early-stage location searches.

EXECUTIVE SUMMARY

- Market context establishing why AI governance — not tool selection — is the first decision an economic development organization has to make, and why the exposure already exists whether or not the organization has acted
- A working vocabulary for AI vendor claims — the difference between retention, training use, and deletion rights, and why conflating them is the most consequential mistake in tool diligence
- A practical guide to the four workflows where AI returns real time in an EDO, including the data conditions each one requires and the reason most implementations fail before the software is opened
- A twelve-question AI diligence framework for executive directors, with guidance on what a credible vendor answer sounds like and what a dodge sounds like
- A machine-readability playbook for EDOs — how to find out what AI systems currently say about your community, and how to fix the specific reasons your data cannot be read

OPENING: THE DECISION NOBODY MADE

Nearly every economic development organization in the country is already using artificial intelligence. Very few have decided to.

I have not yet said that sentence to a room of executive directors without watching several of them go quiet. Not because it is a surprise — most of them suspect it — but because it names the thing everyone has been working around. AI did not arrive at your organization through a procurement process, a board vote, or a budget line. It arrived through a project manager with a deadline and a personal login.

For most of the last three years, the professional conversation about AI in this industry has been a tool conversation. Which platform, which demo, which conference session. That framing has it backwards, and the error is expensive. The tool question is genuinely secondary, and organizations that answer it first tend to end up with software nobody opens and a policy written to retroactively justify a purchase already made.

The reason is structural. An economic development organization is not a private consultancy. It handles code-named projects under confidentiality agreements, it generates records that may be subject to public disclosure, and in most cases it spends public money under procurement rules. Those three obligations collide in ways a general-purpose AI policy template does not contemplate — and they collide the first time a staff member pastes a project description into a consumer chatbot, which for most organizations reading this has already happened.

The organization that treats AI as a technology decision will spend eighteen months discovering it was a governance decision. The organization that markets acreage machines cannot read will lose projects it never knew it was competing for. Both failures are avoidable, and both are avoidable in week one rather than month twelve.

This paper is the framework I use to avoid them. It is written for both sides of the question — the internal governance problem and the external visibility problem — because in practice they are the same problem asked from opposite ends: can the information your organization holds be trusted, and can it be read?

SECTION 1 — HOW AI ARRIVED WITHOUT A DECISION


The Friday Afternoon Problem

An RFI lands at 4:40 on a Friday with a Tuesday deadline. It asks for eleven data points across four sites, a workforce narrative, and an incentive estimate. Your project manager has a soccer game at six.

So she opens a consumer AI account — her own, on her own login, because the organization does not have one — and pastes in the project description to get a first draft moving. The code name goes in. The square footage goes in. The utility load, the timeline, and the incentive range go in.

She is not being reckless. She is being responsive, which is what you hired her for and what her performance review rewards. The organization never told her not to, never gave her a sanctioned alternative, and never explained why it might matter. Multiply that by every staff member with a deadline and you have an unmanaged data flow that leadership cannot describe, inventory, or defend.

WHERE THE PROJECT DATA ACTUALLY GOES



None of this requires bad intent. It requires a deadline.

Figure 1. The path project data takes when no policy exists.

Why the Absence of a Policy Is Itself a Position

Organizations frequently tell me they have not addressed AI yet, as though the matter were pending. It is not pending. An organization with no policy has adopted a policy of unrestricted individual discretion, applied inconsistently across staff, with no record of what was sent where. That position was never approved by anyone, and it is the position the organization would have to defend if asked.

You already have an AI policy. It is whatever your least cautious staff member does at 4:40 on a Friday. The only question is whether you wrote it.

SECTION 2 — ANATOMY OF AN AI EXPOSURE

2A: The Three Obligations That Collide

A private firm that leaks its own confidential data has a business problem. An economic development organization has something more complicated, because three distinct obligations converge on the same act.

Prospect Confidentiality

You sign non-disclosure agreements. Your staff member's personal AI account did not. Terms vary widely on retention, human review, and whether inputs inform future model training — and those terms change without anyone on your team being notified. When you promised a site selection consultant that Project Cardinal would stay confidential, you made that promise on behalf of systems you had never inventoried.

Public Records Obligations

Depending on your structure and your state, prompts and outputs created in the course of official business may be public records. That cuts in two uncomfortable directions at once. If they are records, can you produce them in response to a request — when they live in a personal account you do not control? And if they contain confidential project information, can you properly withhold them under the applicable exemption when you cannot locate them?

Procurement and the Appearance of a Decision

When a staff member adopts a tool, the organization has effectively selected a vendor without a procurement process, a contract, or a data agreement. For an organization receiving public funds, that is a defensible-process problem waiting for an inconvenient moment — usually a moment chosen by someone else.

This paper is general information from a practice that works with economic development organizations. It is not legal advice, and records, confidentiality, and procurement obligations vary meaningfully by state and by organizational structure. Confirm your specific obligations with counsel.

2B: The Vocabulary That Costs Money

An AI vendor claim is a legal and engineering document wearing the clothes of a marketing statement. The terms below are where the real content hides, and they are routinely conflated in sales conversations — sometimes deliberately, more often because the person running the demo does not know the difference either.

- **Retention versus deletion rights:** A vendor that does not train on your inputs may still retain them for months for abuse monitoring. Those are different commitments. Ask for both in writing, with a number of days attached to each.
- **Enterprise terms versus consumer terms:** The same brand often operates under fundamentally different data commitments depending on account tier. A staff member using the consumer version of a tool your organization licenses at the enterprise tier is not covered by the enterprise terms.
- **Training use versus human review:** A commitment not to train on your data is not a commitment that no human will read it. Abuse review, quality sampling, and support access are separate pathways, each requiring a separate answer.
- **Data residency versus data control:** Where the servers sit is a different question from who can compel access to what is on them, and from whether you can extract or delete it on demand.
- **Subprocessors:** Most AI products are assembled from other companies' models and infrastructure. The commitment that matters is the one that flows down to every party in that chain, not the one on the cover of the vendor's security summary.

A vendor security page dated more than twelve months ago is not evidence. In a market where terms, ownership, subprocessors, and underlying model providers all move quarterly, a 2024 policy tells you about 2024.

2C: What a Defensible Position Actually Requires

The good news, and the reason I am consistently optimistic in these conversations, is that the distance from an indefensible position to a defensible one is short and cheap. It is not a software project. It is four artifacts.

- A written policy stating what may and may not be entered into an external AI system, adopted by the board rather than issued informally
- A sanctioned tool with organizational accounts, so staff have a compliant option that is easier to use than the non-compliant one

- A records position, developed with counsel, on how AI-generated material is retained and produced
- A documented staff briefing, which is the artifact that matters when the question is asked eighteen months later

None of that requires a budget amendment. All of it can be completed inside a normal board cycle.

SECTION 3 — WHERE AI ACTUALLY EARNS ITS KEEP

Set the exposure aside, because the opportunity is real — it is simply narrower than the marketing suggests. In the organizations I work with, value concentrates in four workflows, and the ranking matters more than the list.

3A: The Four Workflows

RFI and RFP Response

This is the one that changes an organization's week. The bottleneck is rarely knowledge; it is assembly under time pressure. When site data, workforce narratives, incentive descriptions, and past responses exist as structured, retrievable content, drafting a first response becomes a matter of hours rather than days. The staff member still verifies every figure. She simply stops rebuilding the same narrative from scratch every time.

Prospect Research and Qualification


Synthesizing filings, expansion announcements, supply chain relationships, and news coverage into a usable profile is exactly the work that consumes a business development officer's morning and produces a document nobody reads twice.

BR&E Synthesis

Most organizations conduct dozens of retention visits a year and cannot see across them. Notes sit in a CRM, unread in aggregate. Structured synthesis surfaces the pattern — the third employer this quarter naming the same workforce gap — which is the entire point of the program and the part most often left undone.

Grant Narrative Drafting

Narrative sections restating organizational history, regional context, and program design across applications are well suited to assisted drafting, with substantive and budget sections firmly in human hands.



Illustrative. Your audit produces this map scored against your workflows, your data, and your staff capacity — so the ranking reflects your organization, not a generic industry template.

Figure 2. An illustrative opportunity map. Scoring against your own workflows and capacity will move these positions.

Notice what is not in the high-impact quadrant: predictive targeting models, custom data platforms, and most of what gets demonstrated at conference booths. Those are not bad ideas. They are late-stage ideas being sold to early-stage organizations.

3B: The Constraint Nobody Wants to Hear About

Here is the finding that disappoints clients most reliably: the limiting factor is almost never the tool. It is the condition of the data the tool would need to read.

An AI system cannot accelerate your RFI response if your site inventory lives in a design-laid-out PDF, your incentive documentation is three versions deep in a shared drive, and your CRM holds four spellings of the same employer. It will produce a fluent, confident, wrong answer, and a staff member under deadline pressure will send it.


Bad data plus AI does not produce insight. It produces confident errors at higher volume, with a professional tone attached.

This is why data readiness gets assessed before a single tool is recommended, and why the honest recommendation is sometimes to spend the first quarter on CRM hygiene and a structured site inventory instead. That is a less exciting board update. It is also the difference between a tool used in year two and one quietly abandoned.

3C: Where Organizations Actually Sit

Progress here is a ladder, and skipping rungs does not work. Nearly every organization I speak with is at stage one or two, and the distance to stage three is a decision and an afternoon rather than a budget cycle.

FIVE STAGES OF AI MATURITY IN AN EDO



Most organizations we speak with are at stage one or two. The distance from there to stage three is a policy and an afternoon — not a budget cycle. Stages four and five are where the actual return is, and neither is reachable while stage three is skipped.

Figure 3. Five stages of AI maturity in an economic development organization.

3D: Verifying a Vendor Claim — The Four-Document Test

When a vendor claims a capability, I ask for four things before taking the claim seriously. If a vendor cannot produce them, the demonstration was a marketing artifact.

- The current data processing terms for the specific account tier you would purchase — not the general security page — naming retention periods, training use, human review, and subprocessors
- A written implementation plan tested against your data in its current condition, not against a clean demonstration environment
- Two references from organizations of comparable size and type, with permission to speak to them unaccompanied by a sales representative
- A written statement of exit terms — what you can export, in what format, on what notice, and what copies are retained after termination

This is not an unreasonable ask. Vendors serious about this market produce it in two weeks. The ones that cannot are telling you something important about how the relationship will go once you are committed.

SECTION 4 — THE SECOND FRONT: HOW MACHINES DESCRIBE YOUR COMMUNITY

Everything above concerns what happens inside your organization. A parallel shift is happening outside it, and it is less visible because it produces no notification when it goes badly.

A growing share of early-stage location searches now begin with an AI-assisted query rather than a phone call or a database subscription. A consultant or a corporate real estate team asks a system to assemble candidate locations meeting a set of criteria, and what comes back is a shortlist.

WHERE YOU GET SCREENED OUT

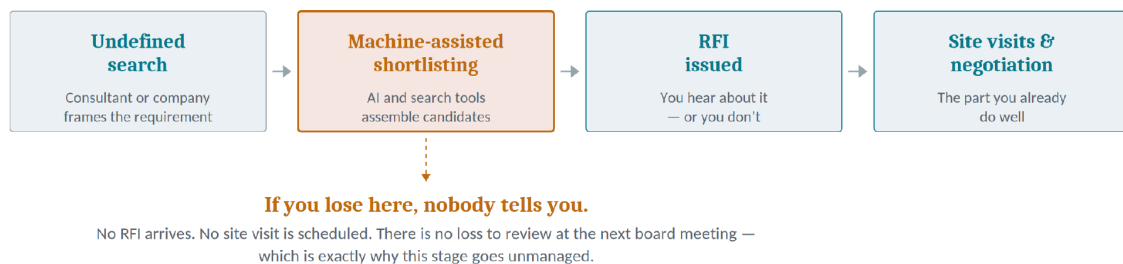



Figure 4. The stage that generates no feedback when you lose it.

If your community is not in that answer, you did not lose the project. You were never considered for it. There is no RFI to review, no debrief, and nothing to report to the board — which is precisely why this stage goes unmanaged.

Why the Data Does Not Make It Through

Most EDO websites are built for a human with time to browse. AI systems do not browse; they parse. And the assets you are proudest of are usually the least parseable things you publish.

THE SAME DATA, TWO OUTCOMES



Trapped in a brochure

A picture of a table. Nothing to extract, compare, or cite.

Site	Acres	Rail	Utilities
North Commerce Park	420	Yes	All 4
Airport Industrial	185	No	All 4
Riverbend Megasite	1,240	Yes	3 of 4
Eastgate Tract	96	No	All 4
County Line West	310	Yes	2 of 4

Structured and readable

Every field labeled, extractable, and citable by name.

Human readers can handle either one. Machines can only use the second.

Figure 5. Identical information. Only one version can be read, compared, or cited.

- **Your best data is trapped in PDFs:** available buildings, incentive summaries, demographic profiles, workforce data, utility rates — all sitting in brochures. A PDF is a picture of a table as far as most systems are concerned.
- **Your site inventory is not structured:** acreage, zoning, utilities, and rail live in a design layout instead of a labeled table, so nothing can extract or compare them.
- **Nothing tells a machine what your entity is:** no schema markup, no canonical fact pages, no consistent naming — so your organization, county, and region get conflated, split, or attributed to a neighbor.
- **You may be blocking the crawlers outright:** a surprising number of EDO sites unintentionally block AI indexers through default content delivery network settings. Nobody chose that; it just happened.
- **Third-party sources contradict you:** state site databases, listing services, mapping platforms, and old news coverage — when these disagree with your website, the machine has no way to know which is right.

The techniques for influencing how AI systems describe a place are new and still being figured out. Anyone promising guaranteed placement in chatbot answers is selling something. What holds up regardless is the underlying work — structured data instead of PDFs, accurate listings, clean markup, current information — which improves traditional search, site selection platforms, and the human consultant hunting your acreage figure at eleven at night.

SECTION 5 — THE AI DILIGENCE FRAMEWORK: TWELVE QUESTIONS

These are the questions I put to a vendor before a tool reaches a shortlist. The value is less in the questions than in listening to how they are answered — precision is a signal, and so is its absence.

Data and Confidentiality

- 1. Where does our data go, which subprocessors touch it, and is any of it used to train or improve models — for the specific account tier we would buy?
- 2. What is retained, for how many days, and can we delete on demand? Put both numbers in writing.
- 3. Under what circumstances does a human at your company or a subprocessor read our inputs?
- 4. If we receive a public records request, can you export every prompt and output associated with our organization, and in what timeframe?

Fit and Adoption

- 5. What specifically does this replace in our workflow, and which staff member stops doing what?
- 6. What condition does our data have to be in for this to work, and have you looked at ours?
- 7. What does the first ninety days require from our staff in hours per week, and what does success look like at day ninety?
- 8. How many organizations of our size and type are using this today, and may we speak with two of them without you present?

Cost, Exit, and Accountability

- 9. What is total cost in year one and year three, including implementation, training, integration, and per-seat growth?
- 10. What is the contract term, and what are the renewal, price escalation, and termination provisions?
- 11. If we leave, what do we get back, in what format, on what notice — and what copies do you retain?
- 12. When the output is wrong and a staff member relies on it, what does your agreement say about responsibility?

A vendor that answers all twelve with dates, documents, and numbers is worth taking seriously. A vendor that answers in adjectives is a risk the organization has not priced.

SECTION 6 — WHAT EDOs SHOULD DO ABOUT IT

Everything above is analysis. Read from the operator's side, it is a to-do list — and it is a list a county EDO with a modest budget can actually complete.

Ask the Question Before You Answer It

Ask your team, without consequence attached, what AI tools they currently use and for what. You cannot govern a practice you have not inventoried, and you will not get an honest inventory if the question arrives with a threat attached. Make it explicitly safe to answer.

Run the Visibility Test and Take Screenshots

Query several AI assistants the way a site selector would — by requirement, not by your community's name — and read what comes back. This is the fastest way to make an abstract issue concrete for a board, and it costs an afternoon.

Issue an Interim Directive This Month

One page from the executive director stating what may not be entered into any external AI tool — project code names, confidential company information, personnel data — pending a full policy. This closes the worst exposure immediately and demonstrates that leadership acted before an incident rather than after one.

Fix the Data Before You Buy the Tool

Inventory what a tool would need to read: site inventory, incentive documentation, workforce data, past RFI responses. Note the format of each. Count how many are PDFs. That count is your real readiness score, and it predicts implementation success better than any product comparison.

Separate the Advice From the Build

The firm that recommends a system should not be the firm that profits from building it. That is not an accusation of bad faith; it is a recognition that a recommendation is unavoidably shaped by what the recommender sells. Independent specification, then competitive procurement, then independent review of delivered work against the written specification.

Write the Specification So a Local Firm Can Bid It

This one is particular to your mission and worth stating plainly. Specifications drafted by enterprise consultants routinely include certifications, service levels, and platform commitments only large national vendors can satisfy — often for no functional reason. Written deliberately, the same requirements can be met by a competent three-person development shop in your own region.

Your mission is growing local employers. An implementation contract awarded inside your own portfolio is a line in your annual report, a reference for that company, and a relationship that outlasts the project. A national vendor gives you a subscription.

SECTION 7 — THREE SCENARIOS

The three scenarios below are illustrative models, not client engagements. They are constructed from the patterns and failure modes described in this paper, with figures set at realistic values for organizations of this size, to show how the framework behaves under pressure. Every number in them is hypothetical. What is not hypothetical is the shape of the problem — I have watched each of these dynamics play out, in some form, on real engagements.

SCENARIO 1: The Policy Written After the Incident

Profile: County EDO, seven staff, \$1.1M operating budget, quasi-public with a public board. Active on a confidential 400-job advanced manufacturing project under a signed NDA with the company's site selection advisor.

Discovery

- A project manager had been drafting RFI responses in a personal consumer AI account for roughly fourteen months, including the code name, site shortlist, and draft incentive structure for the active project
- The organization had no AI policy, no organizational account, and no record of what had been entered
- A local reporter, having heard the project rumor, filed a public records request for communications and work product relating to the code name
- Counsel could not confirm whether the AI account contents were records, could not access them, and could not certify that the response was complete

Solution and Outcome

The organization issued an interim directive within a week, adopted a board policy inside two months, and moved staff to a licensed organizational account with enterprise terms. The records question was resolved on advice of counsel, but the resolution consumed roughly forty hours of executive director time and a legal bill during a live project.

The direct cost was legal fees and lost weeks. The real cost was the conversation with the company's advisor about whether the confidentiality commitment had been honored — a conversation that had no good version.

Key Lesson: The exposure was not created by the records request. It was created fourteen months earlier and simply became visible. A one-page directive at any point in those fourteen months would have prevented the entire episode.

SCENARIO 2: The Tool That Could Not Read the Data

Profile: Regional partnership, fourteen staff, four counties. Purchased an AI-enabled prospect intelligence platform at \$48,000 per year on a three-year term following a strong conference demonstration and board enthusiasm.

Discovery

- The platform required clean, deduplicated company records to be useful; the CRM held roughly 6,000 contacts with an estimated 20% duplication and inconsistent employer naming
- The site inventory the platform was meant to surface existed only as design-laid-out PDFs, so it could not be indexed
- No staff member had been assigned implementation ownership; the tool had been purchased as a capability rather than as a project
- At month eight, three of fourteen staff had logged in during the prior thirty days

Solution and Outcome

The partnership paused the rollout, spent two quarters on CRM cleanup and converting the site inventory to structured records, then relaunched with two named workflow owners and a defined success measure. Adoption reached eleven of fourteen staff within a quarter of relaunch. The organization paid roughly \$64,000 for eight months of a tool it could not use.

The tool was not the problem, and the vendor had not misrepresented anything. Nobody had asked the question in the four-document test: what condition does our data have to be in for this to work?

Key Lesson: Most failed AI implementations at organizations this size are data failures wearing a software costume. The diligence that would have caught it costs a fraction of one year of the subscription.

SCENARIO 3: The County That Was Not in the Answer

Profile: City and county partnership, population 41,000, marketing a 240-acre certified site with rail access and an unusually strong utility position. Consistent RFP participation, but no large project shortlist appearance in four years.

Discovery

- Queried against realistic site selection criteria, four major AI assistants named neighboring counties and omitted this one entirely in eleven of twelve test queries
- The site's specifications existed only inside a PDF brochure linked from a page with no descriptive text
- The state site database listed acreage that contradicted the organization's own website, and a mapping platform still carried a predecessor organization's name and a disconnected phone number
- The website's content delivery network was blocking several legitimate AI crawlers under a default security configuration nobody had selected

Solution and Outcome

The organization converted the site inventory to structured HTML tables, published canonical fact pages for the site and the community, corrected six third-party listings, and adjusted the crawler configuration. Re-tested at ninety days, the community appeared in seven of the same twelve queries with accurate acreage and utility figures.

Nothing about the community changed. The rail access, the acreage, and the utility position were the same on day ninety as on day one. The only change was whether any of it could be read.

Key Lesson: A four-year shortlist drought is usually blamed on product, incentives, or workforce. Sometimes the product was fine and the description of it was unreadable — and that failure mode generates no feedback at all.

CONCLUSION

The question in front of economic development organizations has been misstated. It is not which AI tool to buy. It is whether the organization can describe, defend, and be found — whether it can account for what its staff already do, whether its data is in a condition to be useful, and whether the systems now screening communities can read what it publishes.

The organizations that handle this period well will not be the ones that adopted the most tools. They will be the ones that took a position early, wrote it down, fixed their data, and bought deliberately when they bought at all.

The New Model Requires Four Shifts:

- **Governance first, not governance fourth.** Adopt a position before adopting tools. The reverse order produces a policy written to justify decisions already made.
- **Workflow fit over feature list.** Every capability claim is three claims. Insist on the one that names which staff member stops doing what.

- **Data readiness as a core competency.** Reading your own data's condition is now as fundamental to this business as reading a Phase I.
- **Machine-readability as the competitive product.** For EDOs, being correctly described by the systems that screen communities is worth more than another brochure, and it is cheaper than a marketing campaign.

You are going to have an AI policy eventually. The only question is whether you write it before or after you need it.

The projects are real and they are coming. The tools are not the question. The question is whether your organization can be trusted with the information it holds — and whether anyone can read it.

ABOUT THE AUTHOR

Devin Hillsdon-Smith | Founder & Principal | Hyphen Strategies, LLC

devin@hyphenstrategiesllc.com | hyphenstrategiesllc.com

Devin Hillsdon-Smith is the founder and principal consultant of Hyphen Strategies, LLC, a boutique site selection and economic development advisory firm. With 15+ years of experience spanning public sector economic development and private sector site selection, Devin has advised over \$8 billion in corporate investments across North America.

Devin serves on the board of directors of the Indiana Economic Development Association (IEDA) and other nonprofit organizations. He is a licensed real estate broker and attorney with deep expertise in utility infrastructure analysis, environmental due diligence, and economic development finance. He maintains proprietary utility infrastructure dashboards tracking substation capacity, circuit headroom, transformer lead times, and interconnection queue positions across the Midwest and Mid-Atlantic.

Hyphen Strategies specializes in corporate site selection (\$50M–\$3B+ projects), certified sites programs, regional utilities studies, data center advisory, and affordable/workforce housing development for economic development organizations and growing companies. The firm's AI advisory practice is independent by design: Hyphen builds no software, holds no vendor partnerships, and accepts no referral compensation.

If you would like the visibility test described in Section 6 run on your own region, we will run it and send you the results. It takes about an hour and it costs nothing.