



The Private Intelligence Network

An owned decision-learning system — bold exploration, tied to outcomes


An organization-owned system that turns evidence into a decision, carries the decision into action and learns from what happened next.

FORETODATA FIELD NOTE | 17 AUGUST 2026

EXECUTIVE CASE

A decision, not another chatbot

Illustrative decision, not a customer case study. A strategy lead has one funded investigation to place and a dozen plausible problems competing for it. The Discovery Engine assembles interviews, operating records, public signals, contrary evidence, permissions and missing-data warnings into a ranked shortlist. The lead selects one problem; the evidence, rationale, approval, action and eventual result remain connected.



That is the private intelligence network in miniature: an organization-owned system that turns evidence into a decision, carries the decision into action and learns from the observed outcome. Models, applications and interfaces can change; the governed evidence and decision record remain under organizational control.

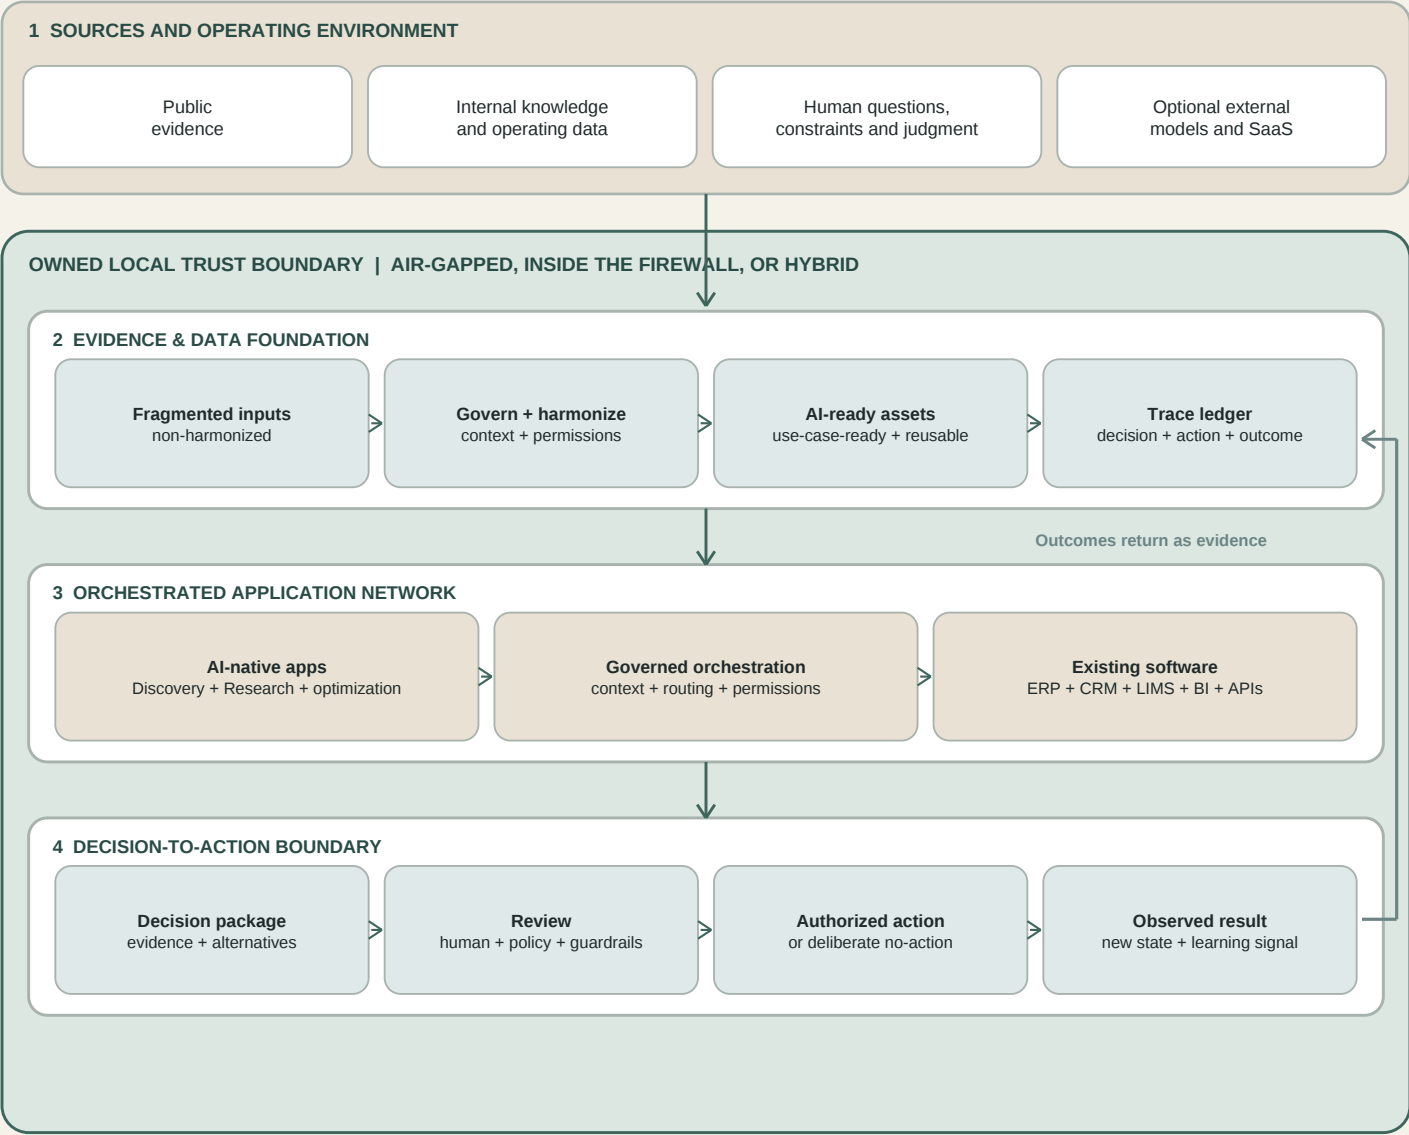
What exists today

- **Observed.** A working local demonstration environment shows evidence-backed discovery, owner-specific decision conversations, governed information sharing, bounded refusals and preserved decision context.
- **Observed local corpus.** A working public-record layer contains 444,248 U.S. patent records—237,261 grants and 206,987 published applications—with ranked semantic search, running on one workstation.
- **Observed scope.** These are working internal demonstrations and sales walkthroughs, not a claimed customer production deployment.
- **Next proof.** Run the first bounded Discovery Engine engagement against a buyer's own decision, evidence, constraints and baseline outcome.

BUSINESS FLOW

Evidence becomes valuable when it changes a decision

The value center is not model output. It is the decision-to-action boundary, where analysis changes research, marketing, sales, strategy, operations or the allocation of real resources. The trust boundary is the architecture border itself - not an extra processing step.



The feedback path is a governed record, not an automatic permission to rewrite evidence or expand an application's authority.

EVIDENCE READINESS


Available data is not decision-ready evidence

Formatting a fragment does not make it true. An evidence record makes its source, uncertainty, permission and intended decision use inspectable so another application can reuse it without silently upgrading a fragment into a fact.

Raw fragment	Governed evidence record
"Release review takes too long." A meeting-note fragment with no defined owner or reuse path.	Claim separated from root-cause proof; linked source records and excerpt; provenance and transformation history; corroboration grade; QA/process permissions; defined decision use.

Foundation first. Leverage next. Expansion after proof.

Begin with one consequential decision, a clear owner and an observable outcome. Build only the evidence foundation and connections it needs; then launch the smallest application capable of improving it.



Current sequence: the Discovery Engine produces a ranked list of problems with evidence of demand. After selection and proof, the Research Engine becomes the second act. Follow-on work reuses the same foundation, permissions, connectors and traces.

One intelligence core, many surfaces


Layer	Examples	Design rule
Governed core	Relational data, documents, vectors, events, semantic layers, knowledge graphs, permissions and traces.	Choose the representation that fits the decision; it is not the user experience.
Delivery surfaces	CRM, ERP, LIMS, dashboards, custom web/mobile apps, desktop or operating-system tools, agents and chat.	Surface governed intelligence where people already work.
Durable record	Evidence, transformations, decisions, approvals, actions and outcomes.	A chat log may be retained as interaction evidence, but it is not the canonical state.

The network is primarily under-the-hood infrastructure. Interfaces may change without discarding the governed intelligence, evidence or institutional memory beneath them.

APPLICATION LAYER

Capabilities tied to potential outcomes

This is the executive-level map. The interactive HTML companion expands one family at a time into its applications and potential outcomes; the leave-behind deliberately keeps the whole shape visible on one page.



Family	First useful deliverable
Shared substrate	A claim whose source, permission and confidence travel with it.
Discover and understand	A shortlist of problems someone is already paying to solve.
Learn and predict	A signal early enough to change the plan.
Decide and operationalize	A feasible action whose tradeoffs and approval are visible.

Potential outcomes are propositions to validate against each organization's baseline, constraints and operating measures. They are not customer results or universal ROI claims. See the companion HTML edition for the expandable application map.

OWNERSHIP

Own the core; choose the boundary

Existing software does not need one forced role. It can be orchestrated inside the network, left outside with manual exchange, or connected selectively through a logged gateway. The evidence foundation and decision trace remain stable across all three choices.

Deployment shape	What crosses the boundary	Best fit
Air-gapped	No live network path. Models, updates and evidence enter through controlled offline transfer.	Maximum custody or highly sensitive work.
Inside the firewall	Internal systems connect locally; approved public material may be retrieved through controlled egress.	Natural default for a shared company system.
Hybrid	The foundation and durable records stay local; selected tasks or approved data may reach external services.	Local custody with optional frontier capability.

"Air-gapped" is accurate only when there is no live network path. A machine disconnected for a period is offline, but not necessarily air-gapped.

Why own the core

- **Data and question privacy.** Source material, prompts, intermediate work and durable records can remain on controlled hardware.
- **Durable custody.** The organization owns its hardware, workflows, evidence store and lawful copies of licensed model artifacts.
- **Policy fit.** Legitimate sensitive work can operate under the organization's own access, biosafety, legal and review controls.
- **Connection without dependence.** External models and SaaS remain optional while the system of record and institutional memory stay local.
- **Reuse instead of reinvention.** One evidence foundation, connector layer, permission model and trace can support many applications.
- **Adoption through products.** Power-user discoveries can become reviewed workflows for colleagues who should not have to become prompt engineers.

LOCAL CAPABILITY

A local model no longer means a toy

Qwen3.8-27B is a compact, downloadable 27B-class multimodal model. Its significance is not that it uniformly replaces hosted frontier models. It is that a model small enough for owned workstation-class hardware can enter the same performance conversation on selected practical evaluations.

Evaluation - higher is better	Qwen3.8-27B	Claude Opus 4.6 Max
Computer use - OSWorld-Verified	84.3	72.7
Agentic coding - SWE-bench Pro	61.7	53.4
Long-horizon office work - CoWorkBench	70.7	68.2
Document intelligence - OmniDocBench 1.5	91.1	86.6
Agentic terminal coding - Terminal Bench 2.1	73.0	78.2
Scientific reasoning - GPQA Diamond	89.2	91.3

These are Qwen vendor-reported point estimates, not an independent parity claim. Harnesses, prompts, tools, quantization, context and workflow design can change results. The card compares against Opus 4.6; Anthropic subsequently released Opus 4.8, so this is evidence of local-model capability, not a current frontier leaderboard.

Anthropic's Opus 4.6 announcement lists standard API pricing of **\$5 per million input tokens and \$25 per million output tokens**. Local inference has no per-token model-vendor charge, but it still consumes hardware, integration, support, electricity and time.

Hardware starting points

A 27-billion-parameter model requires about **54 GB at 16-bit, 27 GB at 8-bit or 13.5 GB at 4-bit** for weights before runtime, context and application overhead. These are arithmetic estimates, not proof of usable speed or context on a particular machine.

Starting point	Observed price	What it honestly establishes
Reuse a current 32-64 GB computer	\$0 incremental	Low-cost proof of fit. Quantized 27B-class work may use system memory or partial-GPU offload and must be timed.
Apple Mac Studio, entry configuration	from \$2,499	Compact unified-memory starting point; configure memory for the target model, context and concurrency.
NVIDIA DGX Spark	\$4,699	128 GB coherent unified memory and 4 TB storage in a purpose-built local AI appliance.

Prices were observed on official U.S. vendor pages on 16 August 2026. They are examples, not quotations or total project budgets. Storage, backup, integration, support, tax, security and operating costs are additional. A machine that loads a model has not yet proved the workflow.

DESIGN RULES

Trustworthy networks are designed around decisions

- 1. Roll out in phases.** Scope the first important decision, construct its required foundation and connections, deploy the highest-leverage application, prove the loop and then extend.
- 2. Start with evidence.** Sources and the Evidence & Data Foundation come before applications.
- 3. Do not confuse available data with AI-ready data.** Fragmented inputs must become fit for a bounded use case; organization-wide readiness adds governed reuse.
- 4. Make applications interoperable.** Every application can read traceable evidence and write governed outputs back to the foundation.
- 5. Separate the intelligence core from its surfaces.** Govern evidence, state, permissions and traces beneath the interface so intelligence can appear in existing systems or purpose-built tools.
- 6. Orchestrate without a rigid pipeline.** Applications can route work and share context without forcing every problem through one sequence.
- 7. Carry the decision record into action.** Evidence, rationale, uncertainty, alternatives, constraints and approvals travel with the decision.
- 8. Reconnect action to outcome.** Decision traces, actions and observed results return as learnable records.
- 9. Govern learning.** Outcomes may improve evaluation and routing but must not silently rewrite evidence or expand authority.
- 10. Preserve deployment choice.** The owned core can run air-gapped, inside the firewall or in hybrid mode through logged connections.

What the learning loop means

The architecture is a **closed-loop decision-learning system**. It becomes literal reinforcement learning only when an application has a valid reward signal, an explicit policy that may change, controlled exploration and evaluation gates. In other cases, the foundation supports organizational learning by preserving which evidence, rationale, decision, approval, action and outcome belonged to the same chain.

A **decision and reasoning trace** does not mean a model's private chain-of-thought. It means auditable decision provenance: sources, transformations, analyses, claims, assumptions, alternatives, uncertainties, rules, tool actions, approvals, overrides, final action and observed outcome.

EVIDENCE NOTES

What is observed, inferred and still to be validated

- **Observed** The architecture can be implemented with an owned evidence layer, local models, existing systems and optional governed external connections.
- **Vendor-reported evidence** The model comparison uses point estimates published in Qwen's model card. It has not been independently reproduced here.
- **Arithmetic inference** The 27B memory figures multiply parameter count by nominal bytes per weight and exclude runtime and context overhead.
- **Potential outcomes** Application outcomes are hypotheses to validate against a buyer's baseline; no universal savings or uplift is claimed.
- **Deployment limits** Air-gapped, firewall and hybrid describe topology choices, not automatic compliance. Controls remain organization- and use-case-specific.

Primary and contextual sources

- **Qwen3.8-27B model card**
<https://huggingface.co/Qwen/Qwen3.8-27B>
- **Anthropic - Introducing Claude Opus 4.6**
<https://www.anthropic.com/news/claude-opus-4-6>
- **Anthropic - Introducing Claude Opus 4.8**
<https://www.anthropic.com/news/claude-opus-4-8>
- **Apple - Mac Studio configuration**
<https://www.apple.com/shop/buy-mac/mac-studio/m4-max-chip-14-core-cpu-32-core-gpu-36gb-memory-512gb-storage>
- **NVIDIA - DGX Spark**
<https://marketplace.nvidia.com/en-us/enterprise/personal-ai-supercomputers/dgx-spark/>
- **Atlassian - identifying AI superusers**
<https://www.atlassian.com/blog/ai-at-work/how-to-identify-ai-superusers>

This field note is an architectural and strategic narrative, not an implementation contract, procurement recommendation, security certification or legal opinion.