VERTICAL-DRIVEN ARCHITECTURES · MINING · DESIGNED WITH PRAXIS · NOVEMBER 2026

Industrial AI Reference Architecture for Tailings Storage Facility Monitoring in the United States

Section Story: Sovereign AI Reference Architectures for Mining at One High-Consequence Facility

Muhammad Umar, Cofounder, CodeNinja

For the accountable executive and the responsible tailings facility engineer, the engineer of record and the independent review board they answer to, and the geotechnical, data and platform engineers who would build and run it.

Vertical-Driven Architectures is a CodeNinja series of system designs. Every design in the series is driven by a real-world problem and scenario in a single industry, and every one is designed on Praxis, CodeNinja's platform for designing physical AI systems. Operations are described by class, never by name.

ABSTRACT

One Section Deserves One Story Before Anyone Asks for It

When an engineer types industrial AI reference architecture for mining into a search box, the question underneath is the one this operation cannot answer today: what is actually happening in one section of a high-consequence tailings embankment in the United States, and is it drifting toward a trigger? The facility holds more readings than ever and the same two engineers reading them, because a piezometer that rises slowly over a month lives in the historian without being joined to the deposition that moved the beach toward it, the rain that fell, the inclinometer two hundred feet away or the radar survey from last quarter. Thresholds in the trigger action response plan fire on single instruments, so operators learn to expect alarms and the pattern that matters hides among them, while the daily inspection app records what the operator saw in words nobody searches. When the engineer of record, the review board or the

regulator asks for the story of one section over a year, it is assembled by hand from six sources, and the assembling takes the very hours the surveillance obligation is meant to spend on judgment.

The design is a sovereign system of context for one high-consequence tailings facility: a governed ontology of 15 objects and their typed links joins 8 named source systems, the historian, satellite radar reports, the inspection app, the trigger action response plan, the engineer of record's records, the deposition and raise plans, the water balance model and the regulator's correspondence, through 2 adapter families, a telemetry adapter and an integration adapter, so that one embankment section has one story. Six services, from the facility ontology and the on-site platform to the phase one gates, serve 5 surfaces, from the trend and trigger view the operators watch to the read-only review view the review board reads. Three drafting agents, a pattern agent, a response agent and a reporting agent, run on gpt-oss-120b, 117B parameters at INT4, 58.4 GB of weights held on one node of eight 80 GB GPUs in the site server room, with anomaly screening on edge units at the logger network and everything on the owner's own hardware. The system improves by memory of confirmed patterns rather than by training, the existing TARP alarm path stays untouched, and the engineer of record holds every interpretation; phase one covers one embankment section for two quarters against the owner's stated measures.

The paper covers the industry problem and the join failure across the existing systems, the constraints and scoping decisions, the layered stack from sources to surfaces, the object model with its typed links, ingestion through adapters, the placement of inference and what never crosses the boundary, the four models and the licenses that let the owner hold their weights, the rollout in three phases with their gates and failure modes, and who owns what the design builds, closing with the chapter on how Praxis contextualized and reasoned the design.

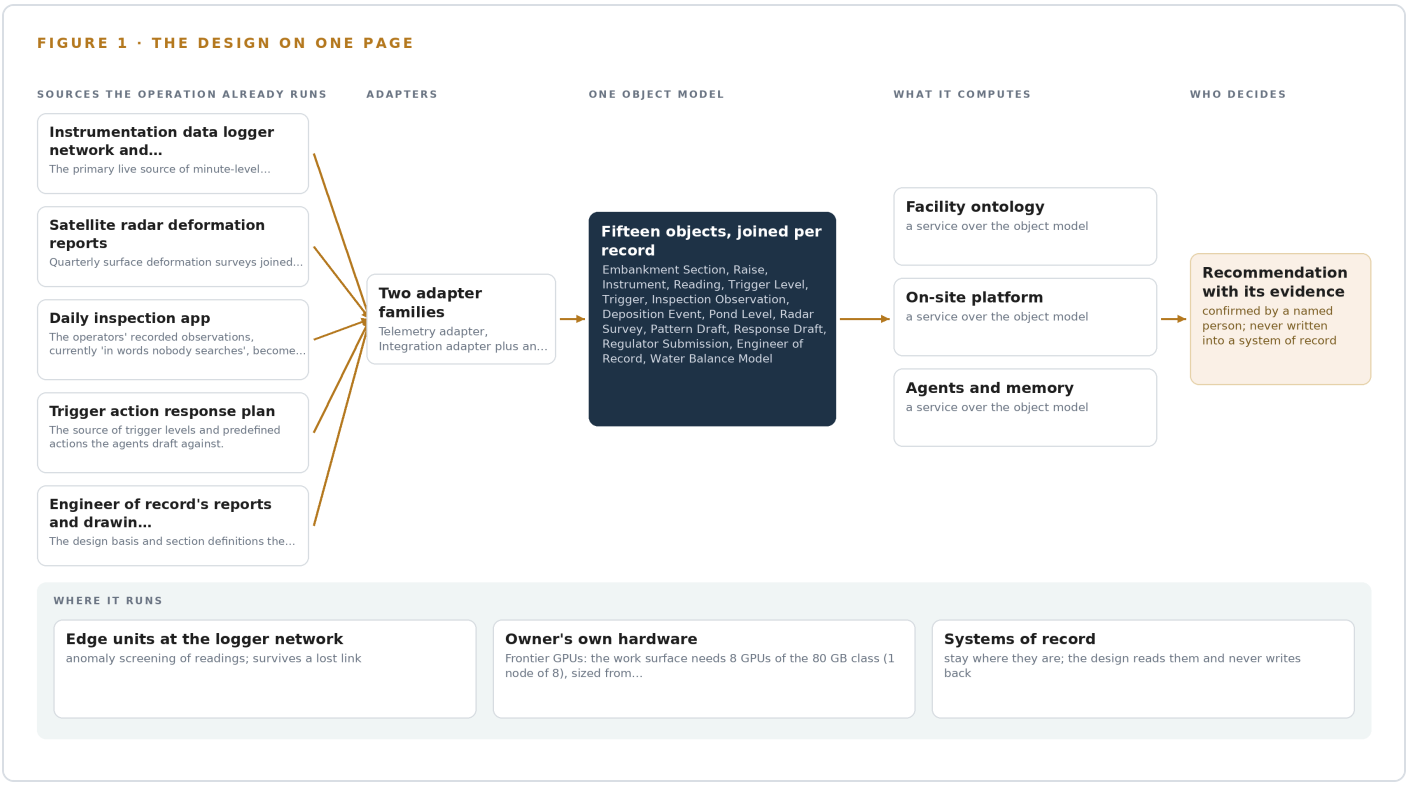


Figure 1. Section Story on one page: the sources the operation already runs, one object model, what it computes, and the person who decides.

Contents

Each chapter is tagged for the reader it serves most directly: Executive, Team Lead, FDE, Reference.

	Abstract · One Section Deserves One Story Before Anyone Asks for It	Executive
PART I · THE PROBLEM		
1	An Industrial AI Reference Architecture for Mining Begins at the Embankment	Executive
2	Every System of Record Sees One Slice of the Facility	Executive · Team Lead
PART II · THE DESIGN		
3	Sovereignty and the Engineer of Record Shape Every Constraint	Team Lead
4	One Stack Runs From Systems of Record to Surfaces	Team Lead · FDE
5	Fifteen Objects Turn Facility Data Into One Argument	FDE
6	Every Source Enters Through an Adapter, Never Directly	FDE
7	Screening Belongs at the Loggers and Reasoning on Site	FDE
8	The License Decides What the Operator Can Own	FDE · Executive
PART III · THE ROLLOUT		
9	Shadow Mode Comes Before Any Draft Is Trusted	Team Lead · Executive
10	The Intelligence Should Stay with the Facility That Produced It	Executive
PART IV · HOW IT WAS DESIGNED		
	Conclusion · The Facility That Keeps Its Own Context Can Answer for It	Executive
11	How Praxis Contextualized and Reasoned This Design	Team Lead · FDE
	Sources	Reference

PART I · CHAPTER 1

An Industrial AI Reference Architecture for Mining Begins at the Embankment

A high-consequence tailings facility under the global standard now holds more readings than ever and the same two engineers reading them, so the pattern that matters hides among single-instrument alarms while the story of one section is still assembled by hand.

An industrial AI reference architecture for mining succeeds or fails at the point where the story of one section has to be assembled, so this chapter starts where the abstract left off: the question the operation cannot answer today, the regimes any answer must satisfy, and the scenario the design has to fit. The abstract promised one embankment section with one story, drafted by named agents on the owner's own hardware; what follows is the ground that promise stands on.

1.1 The Question and the Regimes Behind It

The operation needs one question answered continuously: what is happening in one section of the embankment, read as a whole rather than instrument by instrument. A piezometer that rises slowly over a month is visible in the historian, but the reading only means something joined to the deposition that moved the beach toward it, the rain that fell, the inclinometer two hundred feet away and the radar survey from last quarter. The data required is data the operation already holds, spread across eight named systems of record covering readings, surveys, inspections, thresholds, design basis, plans, water balance and correspondence. Two regimes shape what the answer must satisfy. The global industry standard on tailings management obliges trigger action response plans with predefined actions, prompt notification of the engineer of record when a trigger is crossed, and approval of monitoring reports by the engineer of record together with the responsible tailings facility engineer. State dam safety regulation counts the facility as a water impoundment and expects an operation log and unusual event notices, so every interpretation must trace to a named person and a timestamp.

1.2 The Documented Cost of Not Joining

The industry's own safety reporting shows why this matters beyond convenience. The International Council on Mining and Metals publishes annual fatality and injury data for its member companies, and the pattern in that reporting is consistent: everyday injuries keep falling while fatal and catastrophic events remain a persistent challenge (ICMM n.d.; CIM 2026). Responses across the industry are converging on the same diagnosis, with a national serious injuries and fatalities database being built to turn incident data into insight (CIM 2026b). A tailings facility is where the catastrophic end of that curve lives, and an inquest into a mineworker's death in this industry found record keeping grossly deficient (Nscafoundation 2020). That is the failure mode this design addresses at monitoring scale: records that exist but cannot be joined when the question arrives.

1.3 The Operation as a Scenario

The owner is a mining company in the United States running one large tailings storage facility classified at a high consequence level: an embankment raised in stages over decades, holding tailings behind it with a

beach and a pond, inspected daily and reviewed by an engineer of record and an independent review board. The facility is instrumented with vibrating wire piezometers in the embankment and foundation, inclinometers, survey prisms and monitoring stations, with satellite radar surveys of the surface several times a year. The people in the loop, by role, are the operators who walk the facility daily, the two engineers who read the readings, the engineer of record, the independent review board and the regulator. The physical environments are the embankment crest and faces where the logger cabinets sit, the beach and pond the inspections cross, and the site server room where the reasoning hardware would run, one node of eight 80 GB GPUs. The systems of record number eight, all read-only: the instrumentation data logger network and its historian, the satellite radar deformation reports, the daily inspection app, the trigger action response plan, the engineer of record's reports and drawings, the deposition and raise plans, the water balance model and the regulator's correspondence. Facilities of this class are expected to demonstrate full conformance with the global standard this year, and the industry has already moved from periodic inspection to continuous surveillance, which is the ground the design stands on.

PART I · CHAPTER 2

Every System of Record Sees One Slice of the Facility

Each of the eight existing systems holds one true slice, the readings, the surveys, the walks, the thresholds, the design basis, the plans, the water balance and the correspondence, and none of them can join a piezometer, a deposition, a storm and a radar survey into one story.

Chapter 1 established the question the operation needs answered and the regimes that govern the answer. This chapter shows why no single system can answer it, by walking each of the eight systems of record through what it sees and what it misses.

2.1 What Each System Sees

The historian is the primary live source: it sees minute-level readings for every instrument tag, and it misses which section a tag anchors, which design assumption it verifies and which deposition or storm surrounds any given reading. The satellite radar deformation reports see surface deformation across the facility each survey, and they miss everything between surveys and everything beneath the surface. The daily inspection app sees what the operator saw, in words nobody searches, and misses any join between the walk and the trends running beside it. The trigger action response plan sees thresholds per parameter and the predefined actions attached to them, and misses the pattern that spans instruments, because it fires on single instruments. The engineer of record's reports and drawings see the design basis and the section definitions, and miss live behavior entirely. The deposition and raise plans see intended events with dates and volumes, and miss the instrument response those events caused. The water balance model sees forecast pond elevation from inflows and outflows, and misses the measured divergence read against its context. The regulator's correspondence sees due dates, submissions and

notices, and misses the evidence trail behind each filing. Figure 2 sets these eight slices side by side, each converging on the question none of them can answer alone.



Figure 2. Eight systems, each seeing one part of the answer. The question needs all of them in one place at once.

2.2 What None of Them See Together

What no system holds is the join: one piezometer, one deposition, one storm, one inclinometer and one radar survey, arranged as one story of one section over a year. Today that story is assembled by hand from six sources, and the assembling takes the same two engineers away from reading the readings. Thresholds fire on single instruments, so operators learn to expect alarms and the pattern that matters hides among them. The cost in practice is paid twice, in attention and in traceability: when the engineer of record, the review board or the regulator asks for the story, the answer arrives late and hand-built, and an inquest in this industry has already found record keeping grossly deficient where records cannot speak together (Nscafoundation 2020). The failure is not any one system's failure; it is the space between them, and that space is where the design in Part II puts the object model.

PART II · CHAPTER 3

Sovereignty and the Engineer of Record Shape Every Constraint

Compute stays on the owner's own hardware on site, the existing alarm path stays untouched, nothing is trained into weights, and the engineer of record holds every interpretation, and each of these constraints is hard in this industry for a different reason.

Chapter 2 showed eight true slices and the unoccupied space between them. This chapter states the constraints any answer must respect, because each is hard in this industry for a different reason and each shapes everything downstream, from where the models run to who signs a draft.

3.1 All Compute Stays on the Operator's Own Hardware

The first constraint fixes every inference tier on site: the agentic work surface runs on one node of eight 80 GB GPUs in the site server room, and the anomaly screening of readings runs at the logger network, so nothing crosses the owner's boundary and no third-party model service is used. This is hard in mining because the sovereign posture is an obligation rather than a preference, and it is bound to a named scope item and confirmed with the owner's own compliance owner before design freeze, since no integration answers it by accident. It is also hard physically: a GPU node of this class draws about 10 kW, so the server room's power and cooling must be surveyed before any order, with an in-country colocation room as the fallback if the site cannot host it.

3.2 The Existing Alarm Path Stays Untouched

The single-instrument thresholds in the trigger action response plan keep firing in the historian exactly as they do today, and the design sits beside that path with pattern drafts rather than in front of it with a replacement. This is hard because the alarm path is the one part of the monitoring estate that operators already trust and regulators already recognize: replacing safety-critical alarming would invite scrutiny the design does not need and would put adoption at risk on day one. The cost is a second stream to triage, which the rollout answers with an alert budget per engineer per shift and a false-trigger gate.

3.3 Nothing Is Trained Into Weights

The system improves by memory, not by training: every confirmed pattern and its outcome is kept and read by the next run, and no model is fine tuned on trigger history. This is hard because the industry instinct is to retrain on new data, and because memory only compounds as fast as engineers confirm patterns. The compensation is that every model in the stack carries a permissive Apache-2.0 license, so the owner can hold the weights it runs, inspect them and resynchronize them on its own schedule, with no license trigger tied to use.

3.4 The Engineer of Record Holds Every Interpretation

Agents draft and engineers decide: every pattern draft, response draft and submission carries an approver, and the standard's approval chain, the engineer of record together with the responsible tailings facility engineer, is recorded on each monitoring report before release. This is hard because

accountability under the global standard is named to a person, and the engineer of record's authority over interpretations is a term of a written agreement, not a workflow setting. The design consequence is an instruction gate before anything leaves the boundary, so no draft reaches the regulator unapproved.

3.5 Scoping Decisions

The constraints above force choices about where the work starts and what it touches, and Table 1 records the three that shape the rollout: each buys something specific and pays for it in scope, attention or pace.

Table 1 · Scoping Decisions

DECISION	WHAT IT BUYS	WHAT IT COSTS
Bind one embankment section for two quarters first	Proof on real readings and real inspections before anything scales	The rest of the facility waits for the scale-out phase
Sit beside the existing alarm path	Operators keep the path they already trust, untouched	Pattern drafts add a second stream the engineers must triage
Improve by memory, not by training	The system gets better without a training run or a labeling pipeline	Improvement is bounded by how fast engineers confirm patterns

3.6 What the Design Chose Against

Each major choice closed an alternative, and Table 2 records where, what was picked and what it displaced, ending with the boundary of the work itself: field instruments, hardware supply, historic data migration and round the clock support stay out, because instrument supply, servers and bulk cleansing are the owner's, with evidence support only.

Table 2 · What the Design Chose Against

WHERE	WHAT WAS PICKED	INSTEAD OF, AND WHY
Data foundation	A governed ontology over the eight systems of record	A new central historian or data lake migration, because the ask is to read and join what exists, not replace it
Learning	Memory of confirmed patterns and their outcomes	Fine tuning on trigger history, because the ask rules out training a model on the fly
Reasoning compute	One node of eight 80 GB GPUs on site, holding the 117B-parameter work surface model at INT4	A frontier cluster or a hosted model service, because nothing should cross the owner's boundary
Sensing	The existing instruments and the operators' recorded walks	Cameras and video analytics, because no camera appears anywhere in the ask

WHERE	WHAT WAS PICKED	INSTEAD OF, AND WHY
Scope boundary	No new field instruments, no hardware supply, no historic migration, no round the clock support	A turnkey scope, because instrument supply, servers, cleansing and continuous cover are the owner's, with evidence support only

PART II · CHAPTER 4

One Stack Runs From Systems of Record to Surfaces

Systems of record sit below, never replaced and never modified, one governed object model sits in the middle, and the applications and agents sit above, all on hardware the owner owns.

Chapter 3 fixed the constraints and the scoping they forced: read the systems of record rather than replace them, keep every interpretation with the engineer of record, hold the compute on the owner's own hardware, and leave the existing alarm path untouched. This chapter turns those constraints into a stack whose layers run in a fixed order, from the records the facility already keeps to the surfaces its people already use.

4.1 Three Layers in a Fixed Order

The design follows one pattern. The systems of record sit at the bottom, never replaced, never modified and never loaded directly. One governed object model sits in the middle, a projection over those records rather than a copy of them. The applications and agents sit above it and read only the model. Figure 3 shows the layered stack with its counts: eight source systems below, two adapter families above them, fifteen objects in the middle, four models across three inference tiers, then six services and five surfaces at the top. The pattern fits this problem because the ask is a join, not a rebuild: the owner asked to read the loggers, the historian and the reports and join them, and a new central historian or data lake would spend its budget moving data that already sits where it belongs. The projection rule keeps the middle layer small, since it holds objects and typed links rather than copies of minute-level readings, while every layer runs on hardware the owner owns, from the edge units at the logger network to the inference node in the site server room.

FIGURE 3 · THE LAYERED STACK: 5 SURFACES, 6 SERVICES, 15 OBJECT MODEL, 4 INFERENCE, 2 ADAPTERS, 8 SENSING, 8 SOURCES

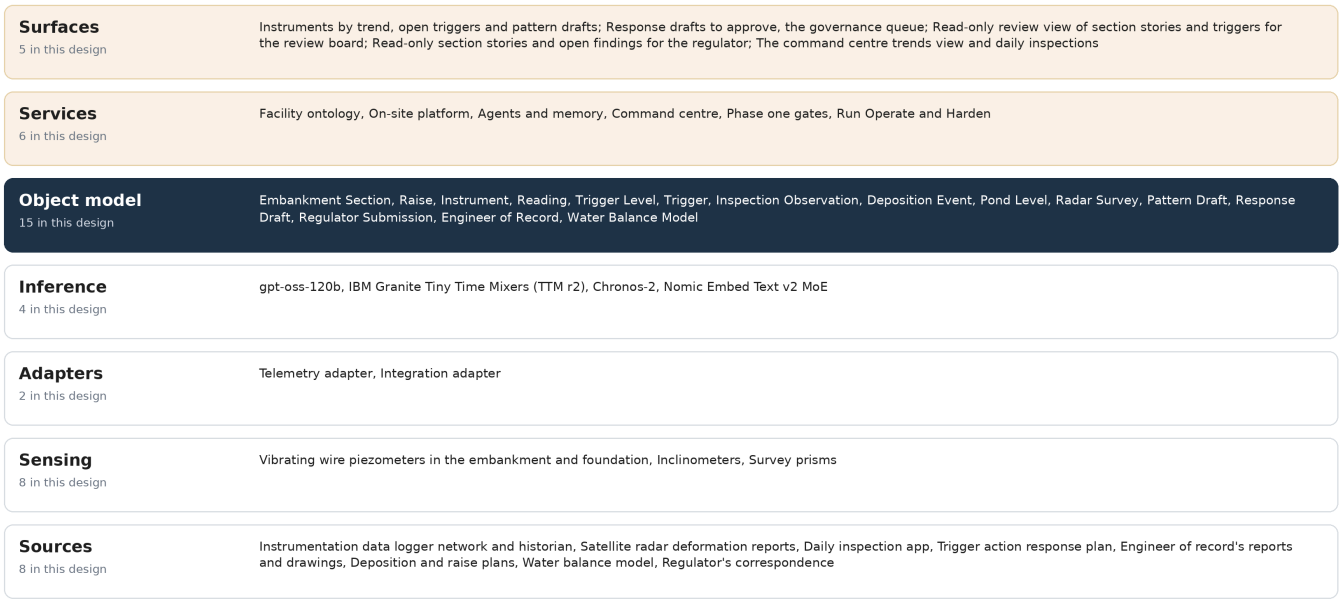


Figure 3. The layered stack: 8 sources, 2 adapter families, 15 objects, 6 services and 5 surfaces.

4.2 The Stack Stage by Stage

Table 3 walks the same stack stage by stage, naming the component responsible at each stage. Two stages deserve comment. The adapters are the single door in: no service, agent or model reads a source system directly, which is what makes the guarantee that no system of record is modified enforceable in code rather than by convention. The inference stage is split by footprint, because a screener small enough to run beside the loggers and a 117 billion parameter reasoning model have nothing to gain from sharing a tier. Table 7, printed with this section by the platform, records the same stack as choices per layer with the reasoning behind each, including the stages this design deliberately leaves out, such as video management and labelling.

Table 3 · The Stack, Stage by Stage

STAGE	WHAT IT IS RESPONSIBLE FOR	HOW
Sources	Holding the records as they stand, in the systems that already own them	Eight systems of record: the instrumentation historian, satellite radar deformation reports, the daily inspection app, the trigger action response plan, the engineer of record's reports and drawings, the deposition and raise plans, the water balance model and the regulator's correspondence
Sensing	Producing the physical readings and observations the records hold	Vibrating wire piezometers in the embankment and foundation, inclinometers, survey prisms, the site monitoring stations, satellite InSAR surveys of the surface, and on-device edge units at the logger network

STAGE	WHAT IT IS RESPONSIBLE FOR	HOW
Adapters	Bringing every source in without touching it	A telemetry adapter for the historian's live feed and an integration adapter for the seven document, plan, model and correspondence sources
Object model	Holding the joined facility as objects and typed links	Hyper Ontology v0.9 carrying fifteen objects, anchored on the engineer of record's section definitions
Inference	Running the four models where their footprint and latency belong	Granite TTM r2 screening readings at the edge on ONNX Runtime under K3s; Chronos-2 and Nomic Embed Text v2 at the site tier; gpt-oss-120b served by vLLM on one node of eight 80 GB GPUs
Services	Doing the work between the model and the surfaces	Six services: the facility ontology, the on-site platform, agents and memory, the command centre, the phase one gates, and Run Operate and Harden
Surfaces	Putting the joined argument in front of the person who decides	Five surfaces, from the trends and open triggers view through the response approval queue to the read-only review and regulator views

Table 7 · Architecture, Layer by Layer

LAYER	COMPONENT	CHOICE	WHY
Event streaming and messaging	by class, no product named	Product	Apache Kafka 4.3.x (KRaft)
The ontology or graph store	by class, no product named	Product	Hyper Ontology v0.9
Time series and telemetry storage	by class, no product named	Product	InfluxDB 3 Core
Video ingest, recording and management	by class, no product named	Not needed	No cameras; inspections are operator walks recorded in the daily inspection app.
Edge model serving	by class, no product named	Product	ONNX Runtime
On site language model serving	by class, no product named	Product	vLLM
Object tracking and trajectory forecasting	by class, no product named	Not needed	No video and no moving-asset tracking in this design.
Edge orchestration of containers	by class, no product named	Product	K3s

LAYER	COMPONENT	CHOICE	WHY
An air gapped artifact and model registry	by class, no product named	Product	Harbor (offline install)
Labelling and retraining	by class, no product named	Not needed	Under system of context nothing is trained into weights; the system improves by memory.
Time synchronisation across sensors and systems	by class, no product named	Product	Linuxptp + chrony
Observability of the running system	by class, no product named	Product	Prometheus 3.13 LTS + Grafana

4.3 A Reusable Pattern

The shape a team carrying a similar brief can reuse is the order itself, not any component in it: sources below, adapters as the single door in, a projection object model in the middle, drafting agents above, and a recorded human approval before anything leaves. Keep the invariants and swap the domain. The fifteen objects become the nouns of whatever facility is at hand; the telemetry adapter's target becomes that site's logger network; the model set changes with the workload, provided the licenses still let the owner hold its own weights. What should never be swapped is the rule that nothing trains on the fly: the system improves by memory of confirmed patterns, and that memory is something the owner keeps.

PART II · CHAPTER 5

Fifteen Objects Turn Facility Data Into One Argument

Fifteen objects and their typed links turn six hand-assembled sources into one queryable argument, so a question about a section reaches its raises, instruments, readings, triggers, inspections, depositions, surveys and submissions across links no document store can follow.

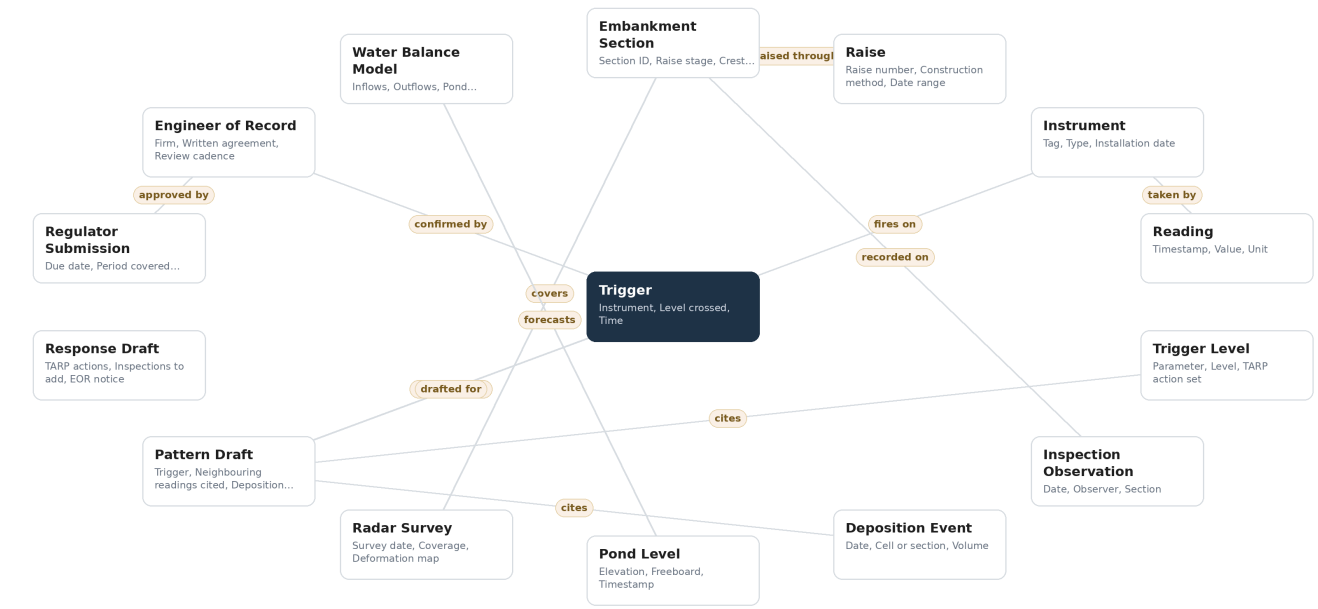
Chapter 4 described the stack that runs from the systems of record to the surfaces. This chapter opens the middle layer, where fifteen objects and their typed links turn six hand-assembled sources into one queryable argument.

5.1 Every Object and Its Typed Links

The model holds fifteen objects. Two are assets: Embankment Section and Instrument. Three are events: Raise, Trigger and Deposition Event. Two are measures: Reading and Pond Level. Four are records: Trigger Level, Inspection Observation, Pattern Draft and Response Draft. Two are documents: Radar

Survey and Regulator Submission. One is an actor, the Engineer of Record, and one is a system, the Water Balance Model. Figure 4 shows every object and the typed links between them. What the links buy is reach. Take the problem the owner stated: a piezometer rising slowly over a month is visible in the historian but joined to nothing. A query anchored on those readings walks to their instrument by tag, to the instrument's section, to the raise stage that section is in, to the deposition events that moved the beach toward it, to the inspection observations recorded on that section, to the radar survey that covered it, and to the matched past episodes held in memory, then returns the whole argument with citations. A document store holding the same six sources can return whatever mentions the instrument tag; it cannot walk relationships nobody wrote down, which is why the story of one section over a year is today assembled by hand. Table 8, printed with this section by the platform, lists every object with the source it is anchored in and its typed links.

FIGURE 4 · THE OBJECT MODEL: 15 OBJECTS AND 12 TYPED LINKS



The navy card is the object the most links reach. A link label is the property the plan carries, or the verb the design names.

Figure 4. The fifteen objects of the model and the typed links that let a query reach across them.

Table 8 · The Object Model

OBJECT	ANCHORED IN	TYPED LINKS
Embankment Section	The engineer of record's reports and drawings	raised through: Raise
Raise	The deposition and raise plans	no typed link recorded
Instrument	The instrumentation data logger network and its historian	no typed link recorded

OBJECT	ANCHORED IN	TYPED LINKS
Reading	The instrumentation data logger network and its historian	taken by: Instrument
Trigger Level	The trigger action response plan with its thresholds	no typed link recorded
Trigger	The trigger action response plan with its thresholds	fires on: Instrument; explained by: Pattern Draft; confirmed by: Engineer of Record
Inspection Observation	The daily inspection app	recorded on: Embankment Section
Deposition Event	The deposition and raise plans	no typed link recorded
Pond Level	The instrumentation data logger network and its historian	no typed link recorded
Radar Survey	The satellite radar deformation reports	covers: Embankment Section
Pattern Draft	record	drafted for: Trigger; cites: Trigger Level; cites: Deposition Event
Response Draft	record	no typed link recorded

5.2 The Human Loop and the Hosting Posture

Every object that carries a decision carries a named person. A Trigger records its confirming engineer; a Pattern Draft moves through Drafted, Confirmed and Rejected only under a human hand; a Response Draft names its approver; a Regulator Submission requires engineer of record approval before it can be Filed. The engineer of record holds authority over interpretations, and the review board and the regulator read the section stories through read-only views. The hosting posture follows the same closed shape. Everything runs on the owner's own hardware on site, so no facility data crosses a boundary. Identity is served on site, so credentials never leave the operation. There are no external links: the radar surveys arrive as report files and the water balance model is read where it runs. The single write path into the object model is the adapter tier, and the sole writes out of it are recorded human decisions, so the systems of record are read-only by construction rather than by policy.

5.3 One Object in Its Recorded Form

The Embankment Section object, printed below in its recorded form, shows the grammar in miniature: identity and location properties carried from the engineer of record's reports and drawings, the raise stage, crest elevation, beach length and consequence classification, and a typed link to the raise it was

built through; the Instrument object is recorded the same way, with a status vocabulary of Commissioned, Degraded, Offline and Retired that a person sets and the agents read.


```
{ "id": "embankment-section", "label": "Embankment Section", "kind": "asset", "anchored_in": "The engineer of record's reports and drawings", "properties": [ "Section ID", "Raise stage", "Crest elevation", "Beach length", "Consequence classification" ], "status_vocabulary": [], "links": [ { "to": "raise", "label": "raised through" } ] }
```


PART II · CHAPTER 6

Every Source Enters Through an Adapter, Never Directly

All eight named source systems enter through one of two adapter families, telemetry for the logger network and integration for everything else, so no source is read directly and no system of record is modified.

Chapter 5 described what the object model holds and who writes to it. This chapter describes how data gets into it: all eight named source systems enter through one of two adapter families, and no source is ever read directly.

6.1 Every Named Source and Its Provenance

Every source in this design is of operator provenance and read-only in direction. The instrumentation data logger network and its historian supplies the minute-level readings for every instrument. The satellite radar deformation reports supply the quarterly surface deformation surveys that join to instrument trends in the section story. The daily inspection app holds the operators' recorded observations, today in words nobody searches. The trigger action response plan supplies the trigger levels and predefined actions the agents draft against. The engineer of record's reports and drawings supply the design basis and the section definitions the ontology anchors on. The deposition and raise plans supply the events that explain instrument behavior. The water balance model supplies the forecast that pond level divergences are read against. The regulator's correspondence supplies the submissions, notices and due dates that are tracked as objects. Figure 5 maps each named system to its provenance class and its adapter path. The industry has reached the same conclusion from the other side: the Global Mining Guidelines Group published a guideline for sharing open data sets in mining because data held in systems that cannot be joined is data wasted (Gmggroup 2022), and China's national mine safety authority issued basic requirements for data fusion and sharing in intelligent mines (Chinamine-safety 2023). This design joins the owner's data inside its own boundary instead of publishing it.

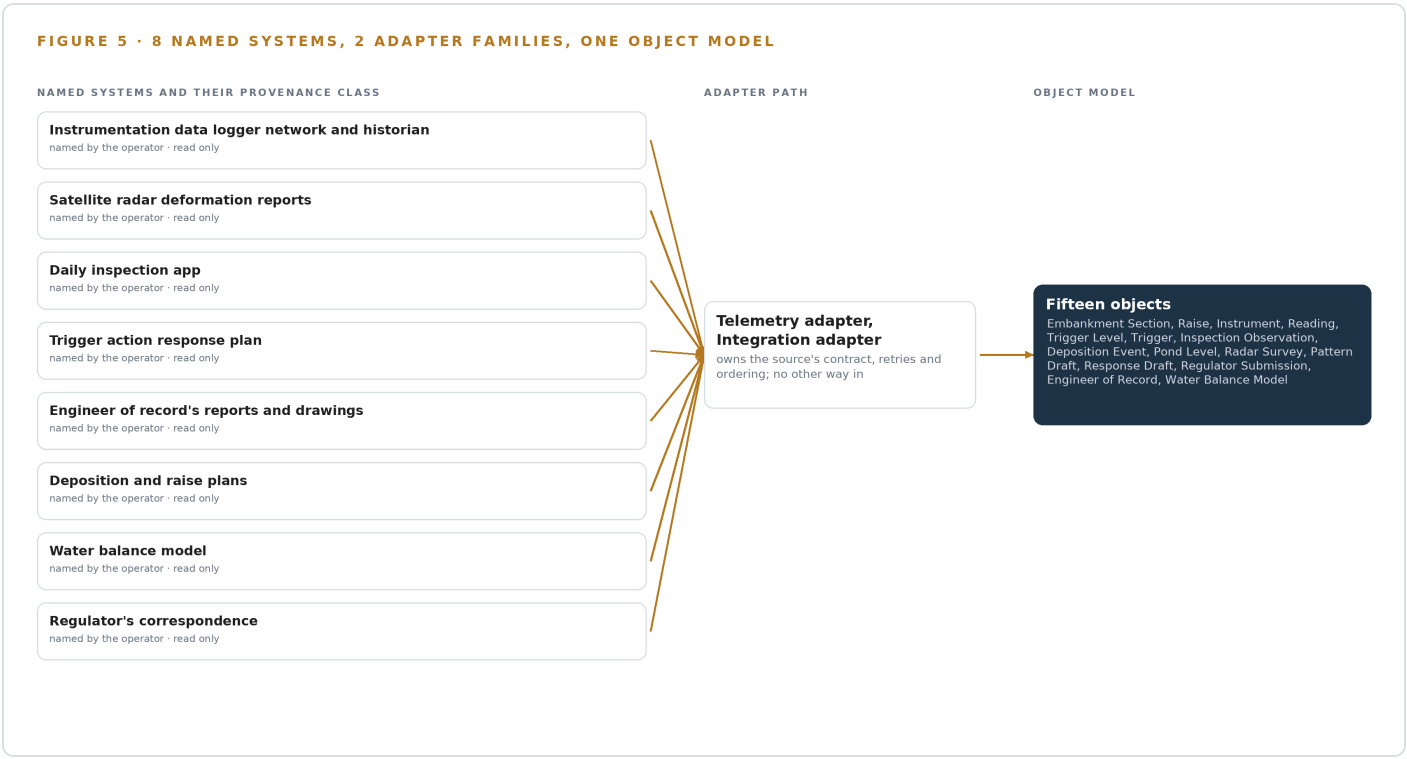


Figure 5. The 8 named systems, the adapter path each one takes, and the object model they all map into.

6.2 What the Adapter Tier Guarantees

The adapter tier makes five guarantees. First, no service, agent or model reads a source system directly; the telemetry adapter and the integration adapter are the sole readers, so the no-modification rule is enforced at one choke point. Second, every source enters with its provenance class and read direction recorded, so any object can be traced to the system that holds it. Third, the telemetry adapter follows the site's own standard, reading the historian's live API or OPC UA feed rather than imposing a protocol, and the integration adapter takes reports, plans and correspondence on a schedule suited to each. Fourth, mappings land in the object grammar of nouns, properties and typed links, and readings carry quality flags so a degraded sensor is visible as such rather than silently trusted. Fifth, the tier applies backpressure, so a surge of readings slows ingestion instead of flooding the model, and no write path ever returns to a system of record, which is what leaves the existing TARP alarm path in the historian untouched.

6.3 The Event Backbone

Between the adapters and the object model sits an event backbone, Apache Kafka 4.3.x in KRaft mode, and its job is to preserve order and causality. Readings are keyed by instrument tag, so every instrument's readings keep their order within a partition even when the fleet of loggers delivers in bursts. Delivery is at least once, with idempotent writes keyed on tag and timestamp, so a replay after a fault cannot double-count a reading. Broker retention buffers the flow when a consumer pauses, and replication inside the site server room keeps the backbone alive through a node loss. Because every causal claim in the section story joins readings on timestamps, time itself is a first-class concern: linuxptp

with an OCP Time Card grandmaster, disciplined by chrony, keeps the loggers, the edge units and the services on one clock. Telemetry lands in InfluxDB 3 Core for trend queries, and every artifact and model update arrives through Harbor in offline install as the air-gapped registry, so even the update path crosses no external boundary.

PART II · CHAPTER 7

Screening Belongs at the Loggers and Reasoning on Site

Anomaly screening runs on edge units at the logger network, trend baselines and the drafting agents run in the site server room on one node of eight 80 GB GPUs, and nothing crosses the network boundary.

Chapter 6 described how every source enters through an adapter and lands in the object model, never directly. This chapter places the compute that reads what lands: anomaly screening on edge units at the logger network, trend baselines and drafting agents in the site server room, and nothing beyond the owner's own hardware.

7.1 Three Tiers and the Memory Arithmetic

The design runs inference on three tiers, and each tier exists because of what it must see and how soon it must see it. The edge tier is a set of on-device edge units placed at the logger network, where Granite Tiny Time Mixers (TTM r2) screens readings for anomalies before they travel anywhere; its footprint is about 0.003 GB at 32 bit, small enough that the precision choice barely matters, and it serves many instrument tags at the gateway on CPU-class hardware under ONNX Runtime. The NPU load on the edge core is unmeasured for this screener, so a dev-board measurement stands as a phase gate before the units are committed.

The site tier runs in the site server room beside the GPU node. Chronos-2, about 0.48 GB at 32 bit as published and about 0.24 GB at 16 bit, builds trend baselines and residual anomalies for cross-instrument patterns, taking rainfall and pond level as covariates so the pattern agent gets context with no training run. Nomic Embed Text v2 MoE, 475 million parameters at about 0.95 GB in fp16, embeds reports and inspection text for retrieval with citations; all of its parameters must stay resident because every expert is live at inference.

The central tier is the drafting tier: gpt-oss-120b served with vLLM on one node of eight 80 GB GPUs, 640 GB of memory in total. The sizing rule the design stands on files 117B parameters at INT4, 0.5 bytes per parameter, which is 58.4 GB of weights; the 1.2 planning factor reserves about 12 GB more for KV cache and activations, so the GPUs must hold 70.1 GB, and the node holds it with 582 GB to spare beside the weights. The model card files the same model as 116.8B total parameters with 5.1B active in a mixture of

experts architecture at MXFP4, about 61 GB, which fits a single 80 GB GPU with headroom. Figure 6 sets the three tiers side by side, with what runs on each and what crosses between them.

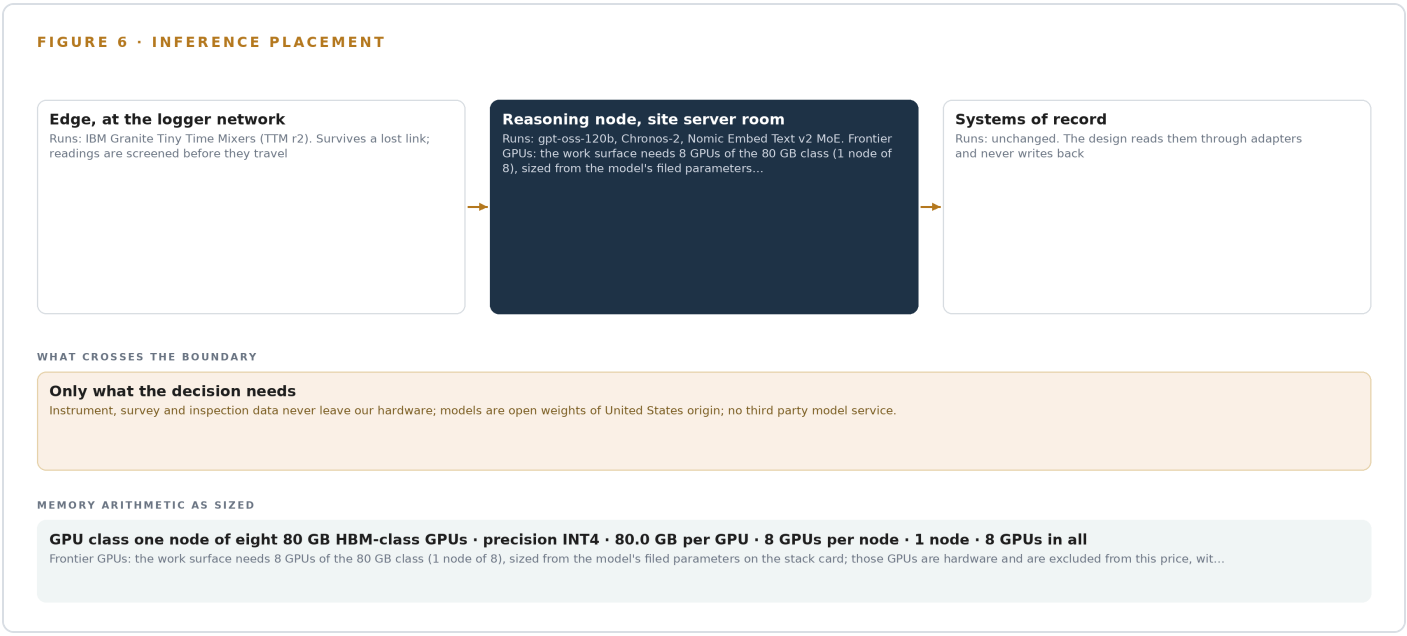


Figure 6. Where each tier runs, what runs there, and the narrow set of outputs that cross the boundary.

7.2 The Latency Budget

The budget is set by the cadence of the work rather than a number on a slide. The historian produces minute-level readings, so edge screening must complete inside the reading interval: a screening result that arrives after the next reading has displaced it screens nothing. Site-tier baselines run on the same cadence, and their outputs feed the pattern agent rather than an operator, so they must be ready when a trigger fires rather than when someone asks. The drafting agents work at the rhythm of an engineering shift, because a pattern draft is useful while the trigger is still open and a response draft is useful while its approver is still on site. The conversational surface must answer with citations at the pace of a conversation, which is where the KV cache ceiling bites: the planning factor's reserve leaves the node's memory mostly free above the weights, so concurrent sessions do not queue behind them.

7.3 The Boundary and the Failure Paths

Nothing crosses the network boundary. All compute sits on the owner's own hardware on site, no third-party model service is called, and if the site segregates its logger network a hardware data diode is the drop-in that keeps it that way. When the link between the logger network and the server room fails, the edge units keep screening against their local model while the event backbone buffers and replays what could not travel, so no reading is lost, only delayed. The existing trigger action response plan alarm path keeps firing on single instruments regardless, because this design sits beside that path rather than replacing it. When power fails, the edge units and the server room restart, the backbone replays its log, and the agents resume from memory, since every confirmed pattern and its outcome is kept and read before the next draft. When the update path fails, Harbor, the air-gapped artifact and model registry

installed offline, holds the newer artifacts and the running versions continue until an update can be carried in; models arrive as artifacts on site, never as a service. One obligation stays open: the on-site compute commitment is bound to a named scope item and confirmed with the owner's own compliance owner before design freeze.

PART II · CHAPTER 8

The License Decides What the Operator Can Own

Four Apache-2.0 models, gpt-oss-120b for the work surface, Granite Tiny Time Mixers for edge screening, Chronos-2 for trend baselines and Nomic Embed Text v2 for retrieval, are all self-hosted, so the owner holds weights it can keep.

Chapter 7 placed screening at the loggers and reasoning in the server room, on one node of eight 80 GB GPUs. This chapter names the four models that run on those tiers and shows why their licenses let the owner hold every weight it runs, which is the question an industrial AI reference architecture exists to answer. Figure 7 sets out the model stack: one model per tier, one license family across all of them.

FIGURE 7 · THE MODEL STACK: 4 MODELS, THEIR PLACEMENT AND THEIR WORK

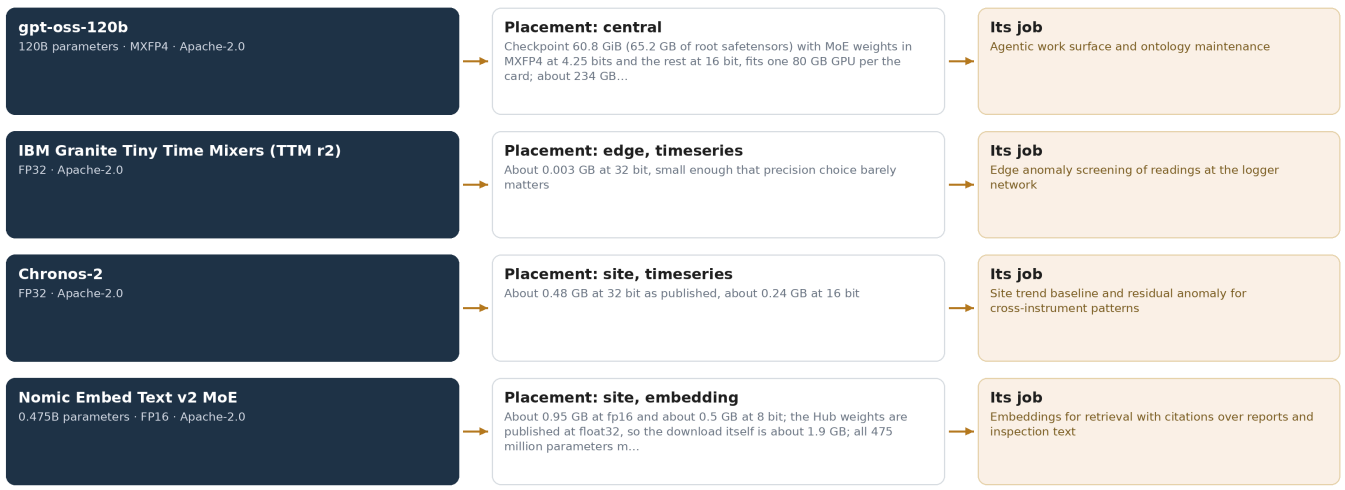


Figure 7. The four models, their placement, and the work each one does.

8.1 Four Models, One License Family

The work surface model is gpt-oss-120b, named by the owner. It is a mixture of experts architecture filed at 116.8B total parameters with 5.1B active, which is why a model of this size fits one node rather than a frontier cluster. Its checkpoint is 60.8 GiB, 65.2 GB of root safetensors, with the MoE weights in MXFP4 at 4.25 bits and the rest at 16 bit; that is about 61 GB, fitting a single 80 GB GPU with headroom, and about

234 GB if expanded to 16 bit. It carries the agentic work surface and ontology maintenance: the named agents that draft, while engineers decide. It is Apache-2.0 and US origin, so the owner hosts it itself with no vendor trigger and no per-token obligation.

The edge screening model is Granite Tiny Time Mixers (TTM r2), an IBM time-series foundation model at about 0.003 GB in 32 bit, small enough that precision barely matters, running under ONNX Runtime on the edge units at the logger network. It screens every reading for anomalies close to the instrument, so only flagged or contextualized data needs explaining downstream. It is Apache-2.0 and US origin; its NPU load on the edge core is unmeasured, so a dev-board measurement is the phase gate.

The trend baseline model is Chronos-2, an Amazon time-series model at about 0.48 GB at 32 bit, about 0.24 GB at 16 bit, running at the site tier beside the GPU node. It gives zero-shot multivariate forecasting with covariates, taking rainfall and pond level alongside instrument readings, so the pattern agent receives trend context without a training run, which the requirement forbids anyway. It is Apache-2.0 and US origin.

The embedding model is Nomic Embed Text v2 MoE, 475 million parameters, about 0.95 GB at fp16 and about 0.5 GB at 8 bit, with the Hub weights published at 32 bit so the download itself is about 1.9 GB. It has an 8K context window and Matryoshka truncation, and its training data is published, which matters for a surface that answers with citations over the engineer of record's reports and the inspection text nobody used to search. It is Apache-2.0 and US origin, and all parameters stay resident because every expert runs at inference.

8.2 The Model and Equipment Register

Table 4 gathers the choices into one register: each model, the hardware classes and sizing rules the footprint rests on, the sensing the models read, the pattern the design stands on, and the ground everything runs on. The register's point is that nothing in it depends on a vendor's continued goodwill: every model is permissive and self-hosted, and the hardware is the owner's own.

Table 4 · Model and Equipment Register

THE CHOICE	WHAT WAS PICKED	WHY HERE
Work surface model	gpt-oss-120b, 116.8B total / 5.1B active, MXFP4, Apache-2.0	Named by the owner; mid size means one node, not a frontier cluster; carries the agents that draft while engineers decide
Edge screening model	Granite Tiny Time Mixers (TTM r2), about 0.003 GB at 32 bit, Apache-2.0	CPU-class footprint serves many tags at the gateway; dev-board measurement gates the unmeasured NPU load
Trend baseline model	Chronos-2, about 0.48 GB at 32 bit, Apache-2.0	Zero-shot multivariate forecasting with rainfall and pond level covariates gives trend context with no training run

THE CHOICE	WHAT WAS PICKED	WHY HERE
Embedding model	Nomic Embed Text v2 MoE, 475M parameters, 8K context, Apache-2.0	Published training data and Matryoshka truncation for a surface that answers with citations
Hardware class	80 GB HBM GPU class	The memory ceiling the reasoning node is sized against
Sizing rule	Frontier model GPU footprint sizing rule (vLLM, SGLang)	117B parameters at INT4 is 58.4 GB of weights; the 1.2 factor reserves about 12 GB for KV cache and activations, so 70.1 GB on 640 GB
Sensing	Vibrating wire piezometers, inclinometers, survey prisms, satellite radar surveys, on-device edge units	Instrument readings and human inspection walks are the sensing; no cameras, and the hardware stays in the United States, so no export control question arises
Pattern	System of context	A governed ontology projected over the systems of record, memory instead of training, named agents that draft and engineers who decide
Ground	The owner's own hardware on site	Nothing crosses the boundary and no third-party model service is called; a hardware data diode is the drop-in if the logger network is segregated

PART III · CHAPTER 9

Shadow Mode Comes Before Any Draft Is Trusted

Three phases carry the work from one section bound into the ontology for two quarters, through operation and hardening, to whole-facility scale-out, with drafts shadowed against the owner's measures before anyone acts on them.

Chapter 8 settled the four models and the licenses that let the owner hold their weights on its own hardware, which is the property that separates an industrial AI reference architecture for mining from hosted analytics. This chapter carries that stack into the field in three phases, each stopped by a gate that can end the work cheaply, with every agent draft shadowed against the owner's own measures before anyone acts on it.

9.1 Three Phases, Fifteen Items and a Gate That Can Stop the Work

Phase One binds one section of the embankment for two quarters of record, and it carries 12 items across five workstreams: the agents and memory, the facility ontology, the on-site platform, the phase one gates, and the command centre. Its exit gate is that one section's records are bound into the ontology, and its item order makes the method visible: bind one section's records, ingest logger readings with edge

validation, then ingest the surveys, inspections and plans that give those readings their context. The Run phase carries 2 items under the on-site platform and the run, operate and harden workstreams, hardening the running platform and then training the tailings team and handing over. Phase Two carries a single item under the facility ontology workstream, scaling the ontology to every section, and it opens only after the Run phase has hardened what Phase One proved. Figure 8 sets the three phases side by side with their item counts, workstreams and gates. Each of the 16 requirements is traced to an item and to a gate, so coverage is a matter of record rather than assertion: an item that answers no requirement is cut, and a requirement with no item becomes an open question rather than a silent gap.

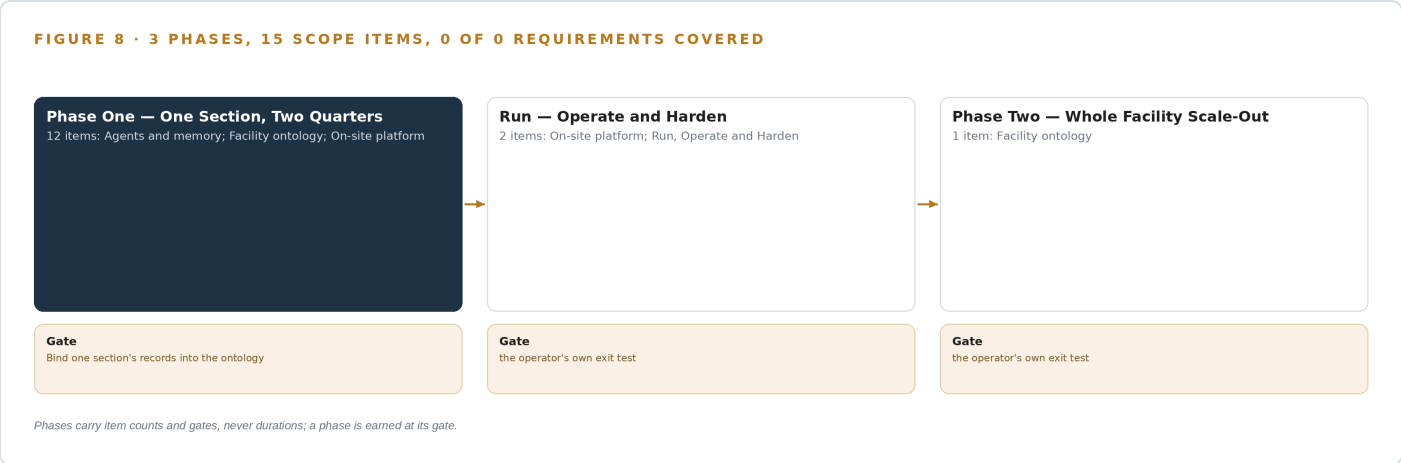


Figure 8. The three phases and their gates, and coverage of the 16 requirements across them.

9.2 What the Rollout Measures

The owner stated four measures, and Phase One is judged against them and nothing else. Around them the design carries gates of its own: a false trigger rate gate on the pattern drafts, an alert budget per engineer per shift so that alarm fatigue is measured rather than argued about, a measurement of the screener's load on the edge core taken on a development board before the edge tier is committed, and data availability checks on the historian at the start of the phase, so that assumed history becomes known history before any memory is built. The sector's own safety reporting shows everyday injuries falling while preventing fatal and catastrophic events remains a persistent challenge (CIM 2026), and the annual fatality and injury reporting across member companies is how the industry watches itself (ICMM n.d.); the measures here therefore weigh the quality of the decision, not the volume of the data.

9.3 Failure Modes and Their Answers

Table 5 states what can fail and what the design does about each, and every row is carried into the phase gates rather than left as a note.

Table 5 · Failure Modes

WHAT FAILS	WHAT THE DESIGN DOES
Alarm fatigue persists if pattern drafts are noisy	A false trigger rate gate and an alert budget per engineer per shift, tuned in the first operating period
The site server room lacks power or cooling for the GPU node	A room survey before any order, with an in-country colocation fallback
An agent-drafted submission is challenged by the regulator	Agents draft and people approve, with approvals recorded under the global standard's reporting clause and an instruction gate before release
The historian's history is shorter or dirtier than assumed	Availability checks at the start of the phase, with assumed rows turned into open questions answered before memory is built
Brittle static liquefaction is not preceded by deformation this sensing can read	The failure mode is carried as unsensed, stated in every document, and governed by conservative design criteria rather than monitoring
The on-site compute obligation is unconfirmed	It is bound to a named scope item before design freeze and confirmed with the owner's own compliance owner

9.4 Lessons

Bind the records before the agents run. The temptation in a facility with more readings than engineers is to reach for the reasoning first, because the reasoning is what the owner asked to see. The design refuses that order: the ontology workstream binds one section's objects and proves the joins on real readings, and only then do the three drafting agents run against a model that can actually answer for itself. Every phase gate in Figure 8 sits on that sequence.

Shadow the drafts before anyone acts on them. A pattern draft, a response draft and a reporting draft are arguments, not actions, and the design keeps them that way until the owner's four measures say the arguments are worth reading. The existing alarm path stays untouched beside them, so the worst case of a noisy draft is a wasted review, never a missed trigger.

State the blind spot in every document. The failure mode that instruments did not see coming in the sector's worst inquiries is carried as unsensed in every story, draft and submission the system produces. A monitoring system that hides its own limits invites the reader to trust it beyond what it can see.

9.5 What Is Still Open

Four questions remain open, and settling each changes something specific. The screener's load on the edge core is unmeasured, and settling it either confirms the edge tier as designed or moves anomaly screening up to the site tier. The server room's power and cooling are unverified, and settling them either places the reasoning node on site or moves it to the in-country colocation fallback, which changes the

boundary. The historian's depth and cleanliness are assumed, and settling them fixes how far back matched past episodes can reach, which bounds what the memory can hold. The on-site compute obligation is not yet bound to a scope item, and settling it with the owner's compliance owner locks the boundary before design freeze, because no integration in the baseline answers it on its own.

PART III · CHAPTER 10

The Intelligence Should Stay with the Facility That Produced It

The object model, the weights, the memory of confirmed patterns and the decision record all stay with the facility that produced them, which is what makes the sovereignty real rather than a hosting choice.

Chapter 9 ended on the open questions the rollout must settle before design freeze. This chapter settles a different question, the one that outlasts every phase: who owns what the work produces, and what that ownership actually consists of.

10.1 What the Operator Holds at the End

The object model is the owner's, and it is a projection rather than a copy: the 15 objects and their typed links describe the facility by joining the eight systems of record, and because the systems of record are never replaced, never modified and never loaded elsewhere, the projection can be rebuilt from them at any time and cannot be held hostage by the design that made it. The weights are the owner's in the strongest available sense: all four models arrive under Apache-2.0, a license with no field of use trigger and no termination on the owner's own decisions, and they run on the owner's own hardware beside its logger network and in its site server room. There are no fine-tunes to own, because the design improves by memory rather than by training, and that is a deliberate ownership choice: an asset that changes only through recorded, reviewable memory cannot drift out from under the people accountable for it. The decision record is the owner's as evidence, not just as data: every confirmed pattern, every approval by the engineer of record and every outcome is kept and read by the next decision, which is what the standard's accountability clauses actually demand. The sector has already seen what deficient record keeping costs in an inquest after a mineworker's death (Nscafoundation 2020), and a record that lives in a vendor's cloud is a record the owner holds on terms it does not set. The boundary, finally, is the owner's: nothing crosses it, identity is served on site, and the single write path into the systems of record is a named person approving a draft.

10.2 The Offer Behind the Design

The design maps directly onto what CodeNinja builds: Adaptive Operations carries the sensing, detection and forecasting of physical behavior, from the piezometer and inclinometer readings through the edge screening and the trend baselines that give the pattern agent its context; Hyper Ontology is the object model itself, the governed 15-object projection over the systems of record that gives one embankment section one story; Hyper Pragma is the work surface where the tailings team runs the named agents that draft while engineers decide, on models the owner can change without losing what was built; Hyper Engram is the kept record of confirmed patterns and outcomes that each run reads before it acts, memory that does not update model weights; Sovereign Infrastructure is the posture the whole design stands on, own hardware, Apache-2.0 weights and operation inside the owner's boundary; and Praxis is the platform that produced this design, which is how every choice in it can be traced to recorded reading.

PART IV · CONCLUSION

The Facility That Keeps Its Own Context Can Answer for It

In one view, the design is a governed projection over eight systems of record the operation already holds: 2 adapter families carry their data into an object model of 15 objects, 6 services and 5 surfaces turn that model into trends, drafts and read-only stories, and 4 permissively licensed models run from the logger network to one eight-GPU node in the site server room, with three drafting agents proposing and named engineers deciding. Nothing is trained, nothing replaces the alarm path, nothing leaves the owner's hardware, and every confirmed pattern becomes memory the next decision reads, so the intelligence the facility produces compounds where the accountability for it already sits, with the engineer of record, the review board and the regulator behind the same story.

Running the same shape elsewhere takes four things rather than a particular industry: systems of record that can be read without being replaced, a named person with the authority to confirm or reject every interpretation, a site able to power and cool one GPU node of the 80 GB class, and model weights under licenses that permit the owner to hold, inspect and keep them. Any operation with instruments, inspections, plans and a regulator's correspondence can bind its own objects through the same grammar, and the first gate, one section bound into the ontology against real readings, is cheap enough to stop the work before it wanders.

PART IV · CHAPTER 11

How Praxis Contextualized and Reasoned This Design

Every choice in this design traces to recorded reading in Praxis, from the ask and its assigned family and industry, through the eight lenses and their cited sources, to the patterns

adopted and the equipment classes the design settles on.

Chapter 10 argued that the intelligence should stay with the facility that produced it. This chapter shows where the design itself came from, because a reader being asked to trust the argument deserves to see how it was reasoned. Every design in this series is produced on Praxis, CodeNinja's platform for designing physical AI systems, and this chapter lets a reader trace any choice back to what justified it (Figure 9). Nothing in the preceding chapters was asserted from preference: each model, each equipment class and each phase gate lands on a recorded read.

11.1 Contextualizing the Ask

The ask was a requirement written in the owner's voice: a system design for monitoring a large tailings storage facility classified at a high consequence level, instrumented under the global industry standard on tailings management and under state dam safety regulation, operated by a mining company in the United States. Praxis assigned the design to the Physical AI family and the Mining industry; the family and industry pin was the designing engineer's own decision rather than a reasoned match, and the brief was still read whole as reasoning input. What Praxis read was the corpus shelf of each lens, assembled before the design was reasoned rather than from searches made to fit the conclusion.

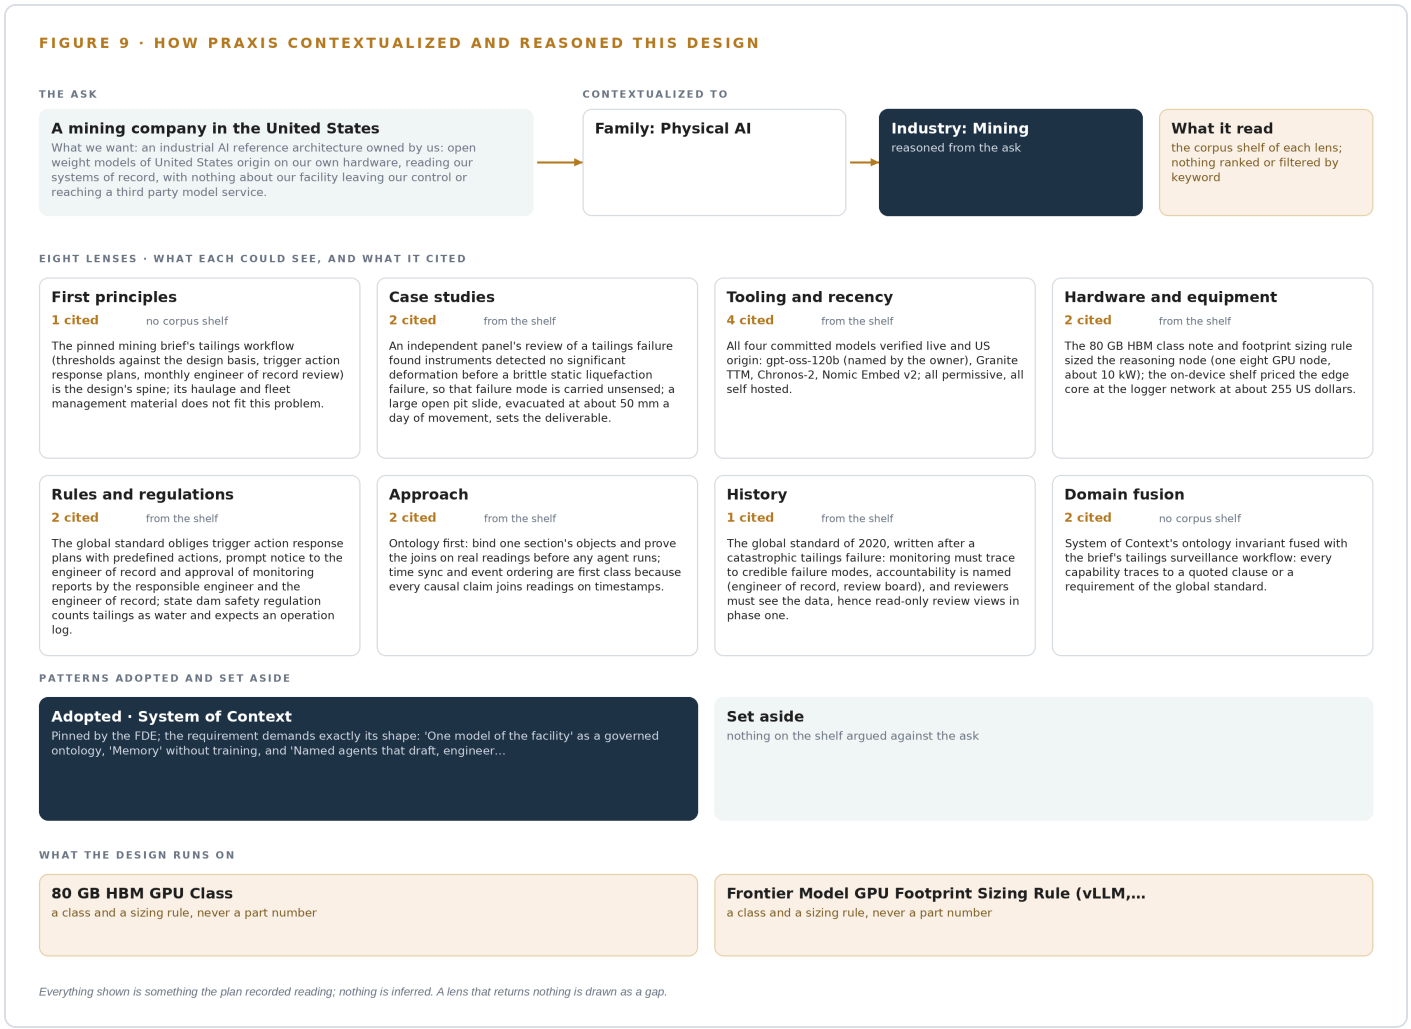


Figure 9. From the ask to the design: the family and industry Praxis assigned, the eight lenses and what each cited, the patterns adopted and set aside, and the equipment the design lands on.

11.2 The Eight Lenses

Table 6 records what each lens could see, how many sources it cited, and what it contributed to the design. No lens returned nothing, but two lenses worked without a shelf, and their contributions came from the brief itself rather than from retrieved records.

Table 6 · The Lenses and What They Contributed

LENS	COULD SEE	CITED	WHAT IT CONTRIBUTED
First principles	no corpus shelf	1	The pinned mining brief's tailings workflow (thresholds against the design basis, trigger action response plans, monthly engineer of record review) is the design's spine; its haulage and fleet management material does not fit this problem.
Case studies	the case study shelf	2	An independent panel's review of a tailings failure found instruments detected no significant deformation before a brittle static liquefaction failure, so that failure mode is carried unsensed; a large open pit slide, evacuated at about 50 mm a day of movement, sets the deliverable.
Tooling and recency	the tooling and recency shelf	4	All four committed models verified live and US origin: gpt-oss-120b (named by the owner), Granite TTM, Chronos-2, Nomic Embed v2; all permissive, all self hosted.
Hardware and equipment	the hardware and equipment shelf	2	The 80 GB HBM class note and footprint sizing rule sized the reasoning node (one eight GPU node, about 10 kW); the on-device shelf priced the edge core at the logger network at about 255 US dollars.
Rules and regulations	the rules and regulations shelf	2	The global standard obliges trigger action response plans with predefined actions, prompt notice to the engineer of record and approval of monitoring reports by the responsible engineer and the engineer of record; state dam safety regulation counts tailings as water and expects an operation log.
Approach	the approach shelf	2	Ontology first: bind one section's objects and prove the joins on real readings before any agent runs; time sync and event ordering are first class because every causal claim joins readings on timestamps.
History	the history shelf	1	The global standard of 2020, written after a catastrophic tailings failure: monitoring must trace to credible failure modes, accountability is named (engineer of record, review board), and reviewers must see the data, hence read-only review views in phase one.

LENS	COULD SEE	CITED	WHAT IT CONTRIBUTED
Domain fusion	no corpus shelf	2	System of Context's ontology invariant fused with the brief's tailings surveillance workflow: every capability traces to a quoted clause or a requirement of the global standard.

11.3 Patterns Adopted and Set Aside

One pattern was adopted, system of context, as the primary, and it was pinned rather than won: the requirement demands exactly its shape, in the owner's own words, one model of the facility as a governed ontology, memory without training, and named agents that draft while engineers decide. The pattern's fixed order governs the stack: systems of record below, never replaced, never modified, never loaded; the owner-owned object model in the middle as a projection rather than a shadow copy, with crossing mappings answering the questions no single system can; one small grammar in which nouns become objects, facts become properties, relationships become typed links and actions change them; and the kinetic loop closing on the model so the sense, decide, act, learn cycle makes it living rather than static. No pattern was set aside, and that absence is itself informative: the ask matched the pattern so exactly that no competing shape survived the first pass, which is why the design spends its rejected alternatives on scoping decisions instead.

11.4 Where the Reasoning Lands

The reasoning lands on two equipment classes that carry the whole footprint: the 80 GB HBM GPU class, which holds the reasoning model's weights with the headroom the memory arithmetic requires, and the frontier model GPU footprint sizing rule for vLLM and SGLang serving, which turned the model's filed parameters into a single eight-GPU node rather than a cluster guess. Everything else in the design, from the telemetry adapter at the logger network to the read-only review views the review board sees, follows from the joins the ontology must make and the boundary the owner set. The closing point is the one this chapter exists to make: everything shown here was recorded reading, the ask, the shelf, the lenses and their citations, and nothing in the design is inferred beyond it.

SOURCES

Source Register

Nscafoundation. 2020. NSCA Foundation. https://www.nscafoundation.org.au/news-item/4764/events-gallery/

ICMM. n.d. ICMM, Transparent Safety Data Reporting. https://www.icmm.com/en-gb/our-work/health-and-safety/safety-data-and-indicators

CIM. 2026. ICMM highlights the need to strengthen critical controls. https://magazine.cim.org/en/news/2026/icmm-highlights-the-need-to-strengthen-critical-controls-en/

Chinamine-safety. 2023. 智能化矿山数据融合共享 数据共享基本要求. https://www.chinamine-safety.gov.cn/xw/zt/2018zt/mkjqryfyy/gzdt_01/202307/W020230704685338069523.pdf

Gmggroup. 2022. GUIDELINE FOR SHARING OPEN DATA SETS IN MINING. https://gmggroup.org/wp-content/uploads/2024/07/2022-04-21_Guideline-for-Sharing-Open-Data-Sets-in-Mining-2.pdf

CIM. 2026. A data-driven safety shift. https://magazine.cim.org/en/news/2026/a-data-driven-safety-shift-en/

About CodeNinja

CodeNinja is a Middle Eastern American artificial intelligence lab and system of context company. It builds sovereign self-improving systems that organizations own, run inside their own boundary and improve with every decision they make.

codeatoms.ai