

CORA



Toward a Connecting Layer for Beamline Experiments

Doğa Gürsoy · Argonne National Laboratory

Fields of View Workshop · MAX IV · May 2026

Humans, instruments, and AI agents reading one record.

A project status, looking for the people who care about the same problem.

Who I am

Mission: close the loop between what we measure and what we know.

FOUR LAYERS, SAME MISSION

Software · TomoPy: open-source CT reconstruction in production at multiple synchrotrons.

Algorithms · iterative methods for imaging with coherent sources.

Scale · distributed reconstruction methods for large-scale CT problems.

Systems · coded-aperture CT prototypes.

Plus 10+ facilities visited as a user, never as a settled local.

What I saw

Every layer, every visit, the same gap.

**This field's knowledge lives in people's heads,
not in anything shared.**

after Polanyi, The Tacit Dimension (1966)

Every visit costs hours just to remember and realign. That is the fragmentation problem, lived. After enough years, you stop accepting that the field cannot have a memory.

Where the context lives

The same kind of information takes a different shape at every phase. The shapes rarely meet.

KIND	BEFORE THE BEAM	AT THE BEAMLINE	AFTER THE BEAM
The sample	recipe, inventory	mounting notes, photos	returned, archived, or destroyed
Live state	scheduled config	EPICS, Tango, Bluesky	rarely retained beyond the run
Records	proposal, safety forms	HDF5 headers, logbook entries	notebooks, paper supplements
Memory	what we agreed	the PI, the student's notes	<i>"ask Dave"</i>

Read across a row. Same kind of thing, three different artifacts, no continuous identity between them.

And what you cannot ask of it

Information is everywhere. Connection is nowhere.

WHAT YOU CANNOT ASK

PIECES THAT EXIST

WHAT'S MISSING

What did we actually run, vs the proposal?

proposal text, live state

nothing **joins** them

Which sample was on the stage during this anomaly?

mounting notes, logbook

nothing **links** sample to event

Everything I ever did with sample A across visits?

sample inventory, per-visit logs

nothing **carries** it across visits

Which raw frames produced figure 3?

TIFF files, paper supplement

nothing **chains** frame to figure

Can an agent steer the next scan?

live state, prior runs

nothing **closes** the loop

Every piece is real. No one owns the connections between them. The picture assembles in someone's head, and only while they are around.

What's been tried

Each layer below solved a piece. None of them owned the connections.

LAYER	EXAMPLES	BLIND TO
Schema	NeXus, CIF/mmCIF, DataCite (FAIR principles overall)	who fills the slots
Catalog	SciCat, ICAT, Tiled	why each output exists
Logbook	Olog, ELOG, eLabFTW	what the instrument was doing
Orchestrator	Bluesky, Sardana, NICOS	anything across visits
Instrument archive	EPICS Archiver, Channel Archiver	the sample, the intent , the decision
Sample tracking	facility LIMS (ISPyB, openBIS), paper notebooks	what the beamline did with it

Read the right column down. That column is the layer **no one** owns.

Why no one has built the connecting layer

Three reasons, in different registers.

Tolerable demand. Humans were the connecting layer. The PI was always around; the student took the notes; the operator remembered the calibration drift. *"Ask Dave"* worked. The absence of machine-queryable connections was not felt as a bug, because the workaround was free at the scale of a single career.

Wrong incentives. Glue between other people's tools does not get funded, papered, or credited (Howison & Herbsleb 2011). Per-facility integrations get built, but on borrowed sysadmin time, and they do not transfer when the sysadmin moves on. The "connecting tissue" role does not exist in any org chart.

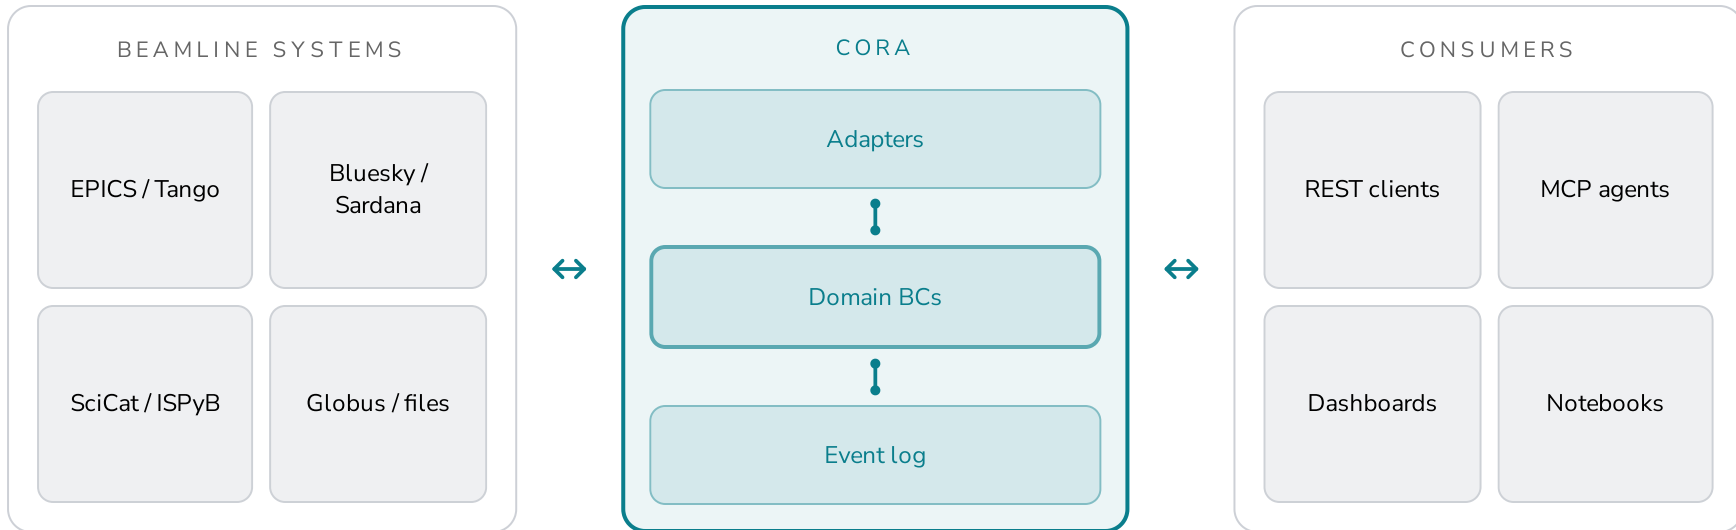
Recent transfer. The patterns are not new: event sourcing began in banking in the 1980s, domain-driven design dates from Evans (2003), Postgres LISTEN/NOTIFY has been production-stable since 2010. What is recent is the cost (cheap enough for a solo project) and the cross-domain transfer into scientific instrumentation (rare until now).

The cultural reason was the load-bearing one. It is about to break: the reasoner is now a model that does not share the room.

What CORA is

So I stopped waiting, and started building the connecting layer.

A queryable causal record of what was set, read, produced, and decided. Shared by humans and AI agents, plugged into existing systems through adapters.



The point is not to record more. It is to make the record one that humans and machines can both reason about.

What CORA is not

CORA is intentionally narrow. Each layer below keeps doing its job.

LAYER	EXAMPLES	KEEPS DOING
Control system	EPCS, Tango	driving the hardware
Orchestrator	Bluesky, Sardana, NICOS	sequencing the scan (<i>CORA fills in where absent</i>)
Data movement	Globus, Tiled, file system	moving the bytes
Analysis	TomoPy, NumPy, PyTorch	doing the math
LIMS / catalog	SciCat, ISPyB	holding samples and finished products
Scheduler	DUO, SMIS	booking the beam
Agent runtime	LangChain, AutoGen, MCP clients	running the agent loop (CORA is the MCP server they call)

CORA reads from and writes to these through adapters. CORA's **job** is the record above them: what happened, who decided, and why. LIMS owns sample logistics; CORA records what the beamline did *with* the sample. Where no orchestrator is in place, CORA can sequence runs itself; where one already drives the beamline, it keeps driving and CORA stays the record.

Why now

The gap was tolerable for decades. What changed: AI as a working participant.

THEN IT WAS TOLERABLE

NOW IT ISN'T

The reasoner was in the room

The reasoner is **a model that isn't**

Joins lived in someone's head

A model only reasons from **what it can read**

AI agents were research demos

Agents **plan, tune mid-run, flag anomalies**

Building this scale needed a team

Humans design, **LLMs implement the scaffolding**

An agent is only as coherent as the record it can read.

AI also lies, drifts, and hallucinates. A unified record is what lets you tell when it is doing those things. CORA does not assume the agent is right; it makes the agent answerable.

A concrete example

In-situ freezing of brine in a porous sample; 30 minutes into a 4-hour scan, an interesting growth front emerges in the streaming reconstruction.

BEAT	WITHOUT A UNIFIED RECORD	WITH CORA
Notice	a glance at the live recon	Run telemetry surfaces the brightening rate and dose-to-date
Recall	<i>"this looked like last winter's NaBr run"</i>	queries prior runs at matched composition
Propose	<i>"let's zoom to 5x there"</i>	Decision proposes lens 1.1x → 5x on the brightest front
Bound	gut check on remaining dose budget	Trust authorizes against the envelope j.park pre-set; Budget gates remaining dose
Act	operator switches the lens by hand	scan-runner agent executes ; Optique Peter ack and new dataset land as events
Audit	<i>"the postdoc remembers when"</i>	one query traces every zoom across the cooling

The cooling happens once. Every zoom either gets recorded against intent, or it gets remembered. Next slide: what that record actually looks like.

What the record looks like

WHAT CORA WRITES

EVENT

```
{
  "event": "ZoomProposed",
  "run_id": "run_2026-05-15_2bm_ice_007",
  "subject_id": "brine_sand_NaBr_5pct_S12",
  "principal": {"kind": "agent", "id": "front_tracker_v3"},
  "from_lens": "1.1x",
  "to_lens": "5x",
  "from_pixel_size": 3.14,
  "to_pixel_size": 0.69,
  "roi_center_px": {"x": 1280, "y": 1024},
  "brightening_pct_min": 8.2,
  "reasoning_ref": "decision_front_pattern_t1820s",
  "ts": "2026-05-15T11:08:42.187Z"
}
```

Immutable. Every actor (human or agent) writes the same shape, signed by the same identity model.

ONE SHAPE

Every event names the principal, the change, and the reasoning ref. Same envelope across all bounded contexts.

ONE LOG

Append-only and immutable at the database level.
Replay rebuilds any projection from scratch.

WHAT YOU CAN ASK BACK

QUERY

```
$ cora why "S12 lens change at t=1820s"

11:08:42.0 RunReading
           brightening 8.2%/min, dose 3.1/8.0 kGy

11:08:42.2 ZoomProposed
           by agent front_tracker_v3
           lens: 1.1x → 5x, ROI (1280, 1024)

11:08:42.3 EnvelopeChecked
           pre-authorized policy (operator: j.park)

11:08:42.4 LensChanged
           by agent scan_runner
           Optique Peter ack: pixel 3.14 → 0.69 μm
```

Four hundred milliseconds, four events. The reason the lens changed is in the record, not in someone's head.

ONE QUERY

"Why did X" returns the causal chain in order. Not a search across files; a walk along the log.

Audit, explainability, and reproducibility become query operations, not human reconstruction.

How it's designed

Six choices. Each is well-known in industry. The new part is the combination, applied to a beamline.

CHOICE

WHAT IT GIVES YOU

Domain-driven design (bounded contexts)

each domain **owns** its rules; one place to change them

Event-sourced storage

every state change is one **immutable** event; replay is free

Functional core, side effects at edges

pure decisions; **testable** without the beam

REST + MCP twin surfaces

every command on both, **same handler**

Agents as principals

humans and LLMs **share** identity, authz, audit

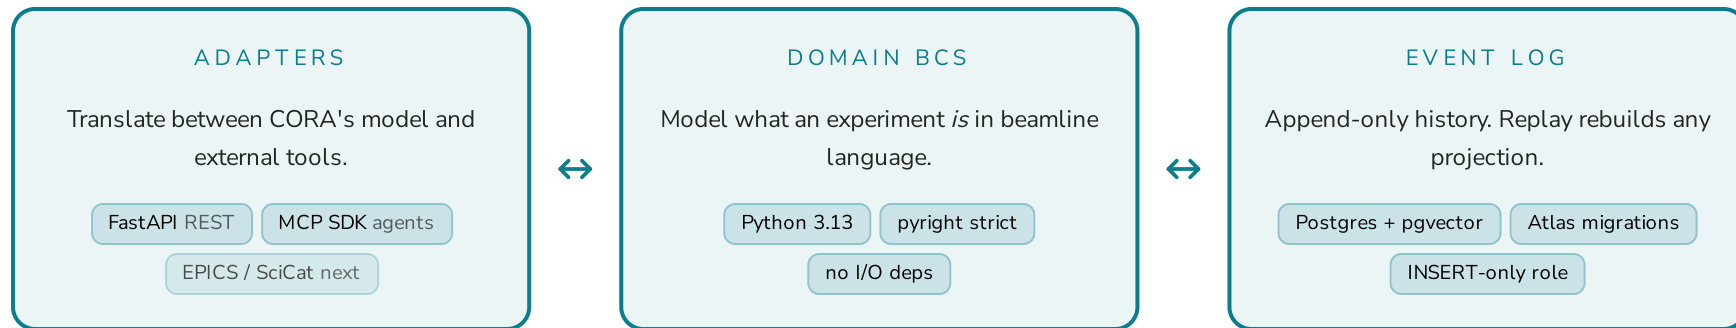
Recipe ladder (Method → Practice → Plan → Run)

abstract intent flows to concrete execution; **facility-neutral**

None of the rows are individually novel. The combination, applied to scientific instrumentation, is.

How it's built

The CORA card, opened up. Three layers, each with a single job.



Today: 14 bounded contexts · 20 aggregates · 9,963 passing tests · pyright 0/0/0 strict · forward-only migrations

The repo has more under each layer: per-BC aggregates, deciders, projections. But this is the picture to leave with.

Where the project is

Architecture stable. Pilot in flight. Honest about both.

BUILT AND STABLE · 14 BOUNDED CONTEXTS

Access · actor identity (register, deactivate)

Trust · zones, conduits, policies

Equipment · family, affordances, settings, ports, condition

Recipe · Capability template · Method → Practice → Plan → Run → Dataset

Subject · sample identity across stations

Run · execution, readings, procedure

Campaign · multi-run studies (in-situ, operando, sweeps)

Data · datasets with lineage and intent

Decision · proposals, approvals, reasoning refs

Supply · reagents and consumables state

Operation · procedure execution (ISA-106 lens)

Safety · clearances, hazard classifications, approval chains

Caution · operator tribal-knowledge cautions and workarounds

Agent · AI agents as principals (RunDebrief, CautionDrafter live)

DESIGNED, NEXT · 2 MORE IN SCOPE

Strategy · *decision modes (human / AI / hybrid) for workflows*

Budget · *allocation tracking (hours, storage, USD, tokens) with limits*

PILOT LADDER

1 **APS 35-BM**
integration in progress

2 **Other APS imaging beamlines**
next

3 **Cross-facility (MAX IV the natural step)**
the conversation I'm here to start

Solo, multi-year horizon. Architecture stable; pilot integration ongoing; not yet handling production beamline traffic.

WHERE I'D MOST WELCOME CHALLENGE

High-frequency telemetry · cross-experiment sagas · agent trust boundaries · schema governance across facilities.

What I'd want from you

Depending on who you are, different things help:

	IF YOU RUN A BEAMLINE	IF YOU BUILD SCIENTIFIC SOFTWARE	IF YOU WORK ON AI FOR SCIENCE
Show me what breaks	what about the model would not fit your facility	which integration points your stack would expose	what context your agents actually need from instruments
Tell me what's missing	how a real shift actually unfolds (I'd love a half-day walkthrough)	the adapter shape you'd want; open an issue	what would close the gap from demo to working tool

`dgursoy@anl.gov` · `github.com/xmap/cora`

Issues, design memos, BC maps: all public. The repo is the work.

Three things to remember

- 1 The connecting layer between instruments, samples, data, and decisions has been missing for decades. AI is what makes it urgent.
- 2 CORA is one honest attempt at that layer: event-sourced, REST + MCP, beamline-vocabulary, open-source. Complementary to NeXus, EPICS, SciCat, not a replacement.
- 3 It exists today: **14 bounded contexts, 9,963 tests, APS 35-BM pilot in progress.** Looking for collaborators who care about the same problem, including this room.



Thank you

The most useful thing for me today is what felt off, or what felt missing.

Do not sugarcoat it.

Q&A