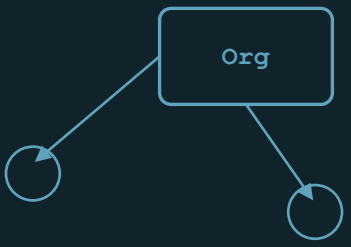

GRAPH DATA COUNCIL · TWENTY-FIRST TUC MEETING · BOSTON

Compliance-Native Graph Schema Design

Encoding FDA 21 CFR Part 11 directly into a property graph — patterns from a production pharmaceutical R&D platform.

Sriram Boppana

Holonyx · September 4, 2026

One product lifecycle. Four disconnected systems.

Every therapeutic is discovered, developed, manufactured and released on separate software — each with its own database and its own audit trail. Compliance is what breaks at the seams between them.

THE PHARMA / BIOTECH LIFECYCLE

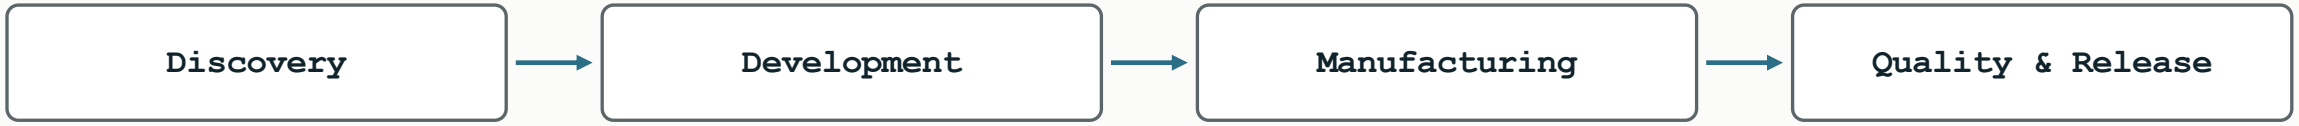

ELN	LIMS	MES	QMS
electronic lab notebook	*sample & test data*	*batch execution on the floor*	*deviations, CAPA, batch release*

THREE COMPLIANCE PROBLEMS TODAY		HOW ONE GRAPH DATABASE ANSWERS
1 Each system keeps its own audit trail, in its own format. Proving nothing was altered means trusting four or five separate logs.	→	**One SHA-256 hash-chained audit trail, inside the graph. Verified by walking the chain — one endpoint, one query.**
2 “What material and which changes led to this batch?” is a manual reconciliation project across systems.	→	**Typed causal edges — DERIVED_FROM, MANUFACTURED_INTO, TRIGGERS. Donor-to-batch provenance is a single traversal.**
3 “Is this process still qualified?” is a periodic manual review — stale the moment it is filed.	→	**Qualification state is a live graph node, recomputed automatically on every batch-record signature event.**

One graph. Typed edges do the work.

Organization is the multi-tenant root — every entity reaches it through one BELONGS_TO edge. Causal edges like DERIVED_FROM and TRIGGERS do what foreign keys across five separate databases never could: they make the relationship a first-class, traversable fact.

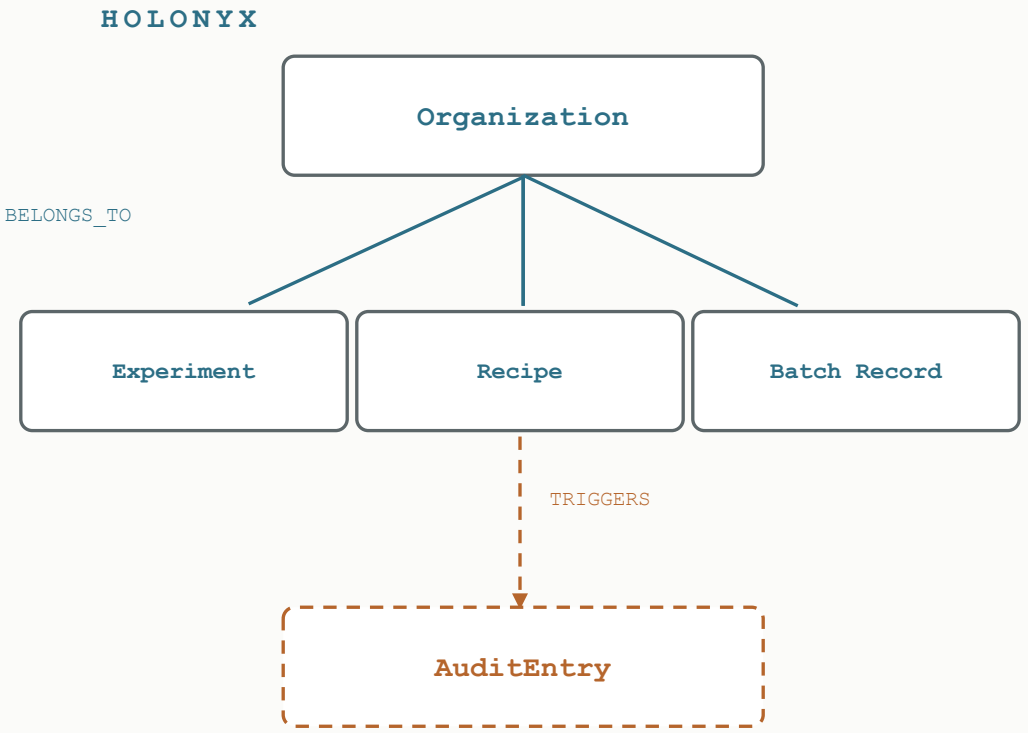

102

DISTINCT NODE TYPES

228

CONSTRAINTS & INDEXES

1

DATABASE — NO SYNC LAYER TO MAINTAIN

```
MATCH (br:BatchRecord {id:$id})
OPTIONAL MATCH p = (dr:DonorRecord)<-[:DERIVED_FROM_DONOR]-(:ApheresisProduct)-[:MANUFACTURED_INTO*1..2]->(br)  RETURN dr, p
```

one traversal covers fresh or cryopreserved starting material — real query, running in production

The audit trail is graph nodes, not a log file

Every state-changing write creates an AuditEntry node whose hash covers the previous entry's hash — a chain you verify by traversal, not by trusting a separate append-only store. The chain lives in the same graph as the data it audits.

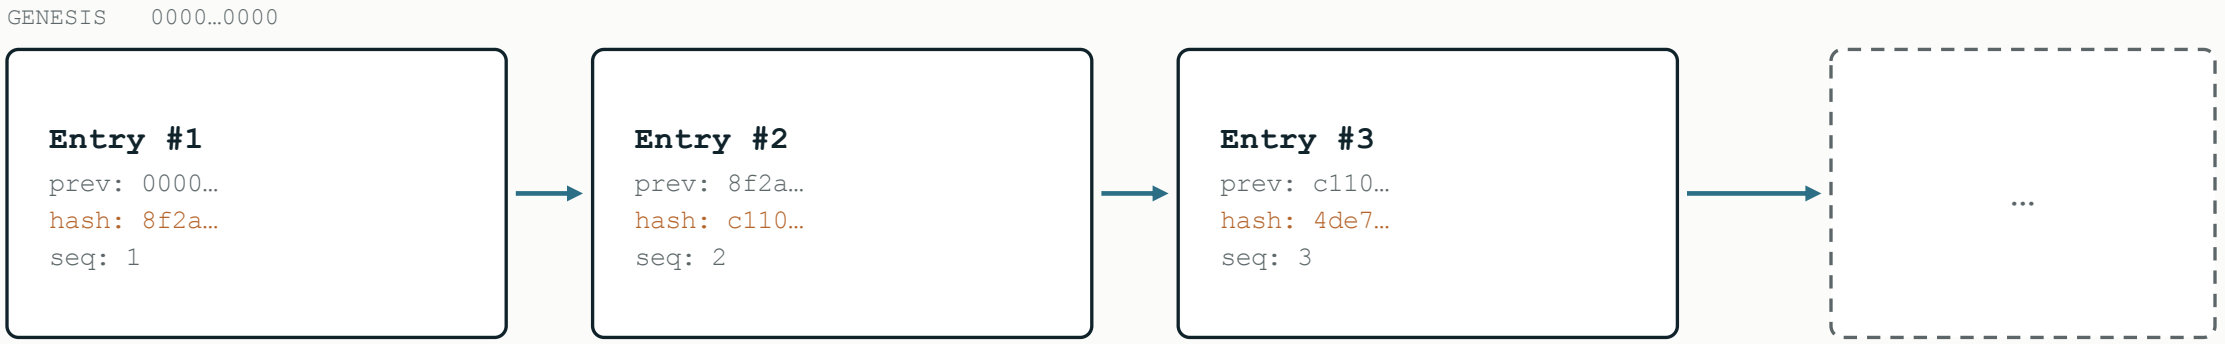


```
chain_input = prev_hash | id | timestamp | userId | action | recordId | newValue   →   SHA-256
```

21 CFR §11.10(e)

verify-chain walks every AuditEntry in sequence order and reports the exact sequence number where the chain breaks.

The graph is the assistant's working memory

At session start the server queries the graph directly and assembles a compact context block — no vector store, no stale snapshot. Refreshed every session:

1. **Recent experiments**

   five most recently updated Experiment nodes

2. **Batches with drift**

   BatchRecord joined to its QualificationState

3. **Inventory expiring**

   InventoryLot nodes inside a 30-day window

4. **Pending approvals**

   AgentTask nodes awaiting a human decision

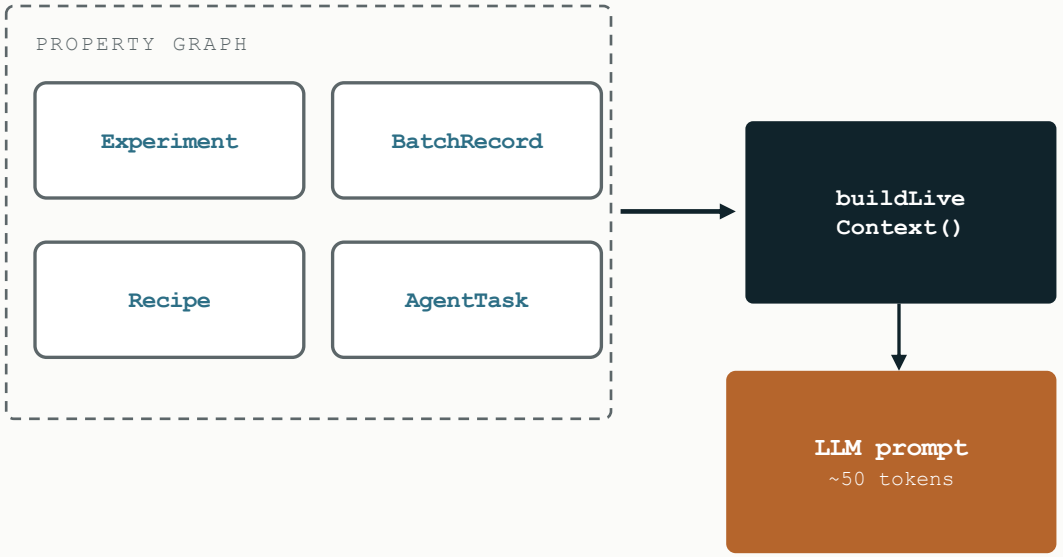

Role-gated write-back

Every AI tool call passes the same role-permission check a human's request does — and every authorized execution writes the same AuditEntry hash-chain entry. The assistant doesn't get a quieter audit trail than a person.

5 PATTERN THREE

Quality state is a node, kept current by real events

Every batch-record signature re-evaluates out-of-spec rate against the last 20 runs and writes the result to a QualificationState node — replacing a periodic manual review with an always-current, queryable fact.

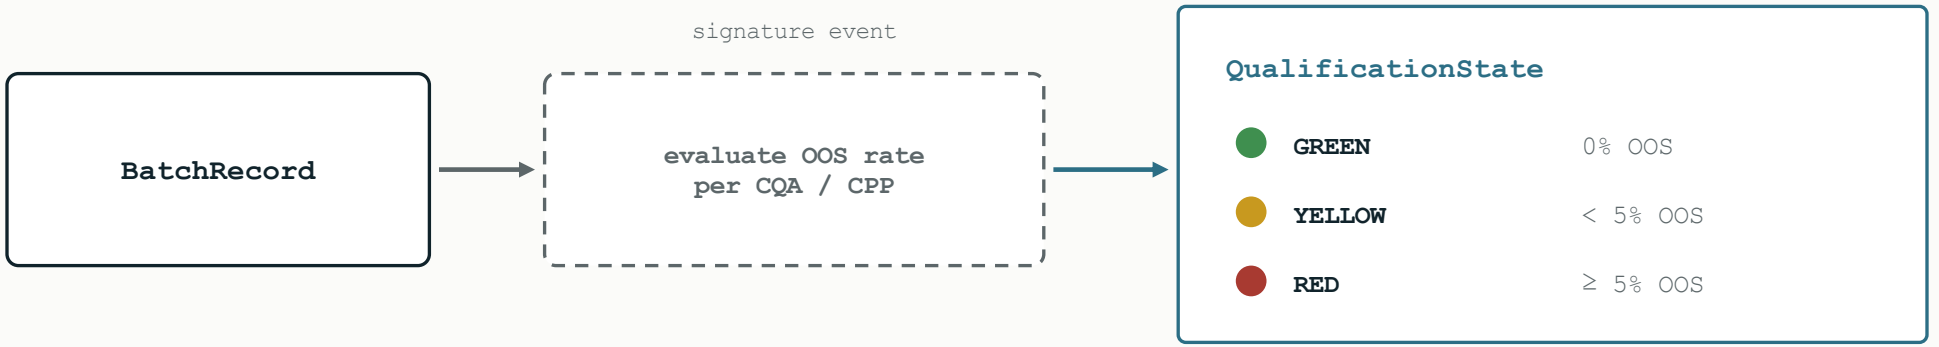

Real-time

NOT PERIODIC MANUAL REVIEW

20

MOST RECENT BATCH RUNS EVALUATED

Why this matters beyond pharma

Any domain with a periodic-review compliance pattern — SLA monitoring, safety-case review, model-drift monitoring for other ML systems — can move that review from a calendar trigger to a graph-state trigger, and get an audit-ready history of every transition for free — it is just more AuditEntry nodes on the same chain.

Holonyx

Thank you

Srirama Boppana

sri@holonyx.io