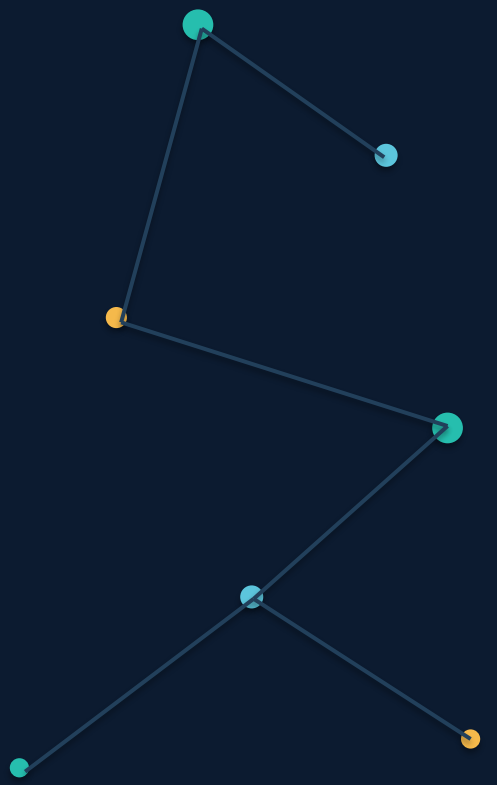
AI IN HEALTHCARE SUMMIT | EXECUTIVE PERSPECTIVE

Building the Foundations for Agentic AI in Healthcare



From multimodal data to context-aware, governed intelligence

DATA • CONTEXT • GOVERNANCE

Speaker: Nirav Parikh, Data and AI Executive

The shift is from generating answers to coordinating work

Agentic AI changes the unit of value, from a response to an outcome.

GENERATIVE AI

Creates and summarizes content

Responds to a prompt

Human coordinates the next step

AGENTIC AI

Interprets a goal

Plans and uses approved tools

Coordinates multi-step workflows

HEALTHCARE REALITY

High consequence decisions

Fragmented clinical context

Complex people + system workflows

“The next frontier of healthcare AI is not a better model; it is a better foundation for intelligence.”

Healthcare information is rich, but fragmented

The challenge is not just data volume. It is meaning, connectivity, timing, and trust.

EHR + encounters

Imaging + DICOM

Genomics + claims

Notes + documents

Pathology + labs

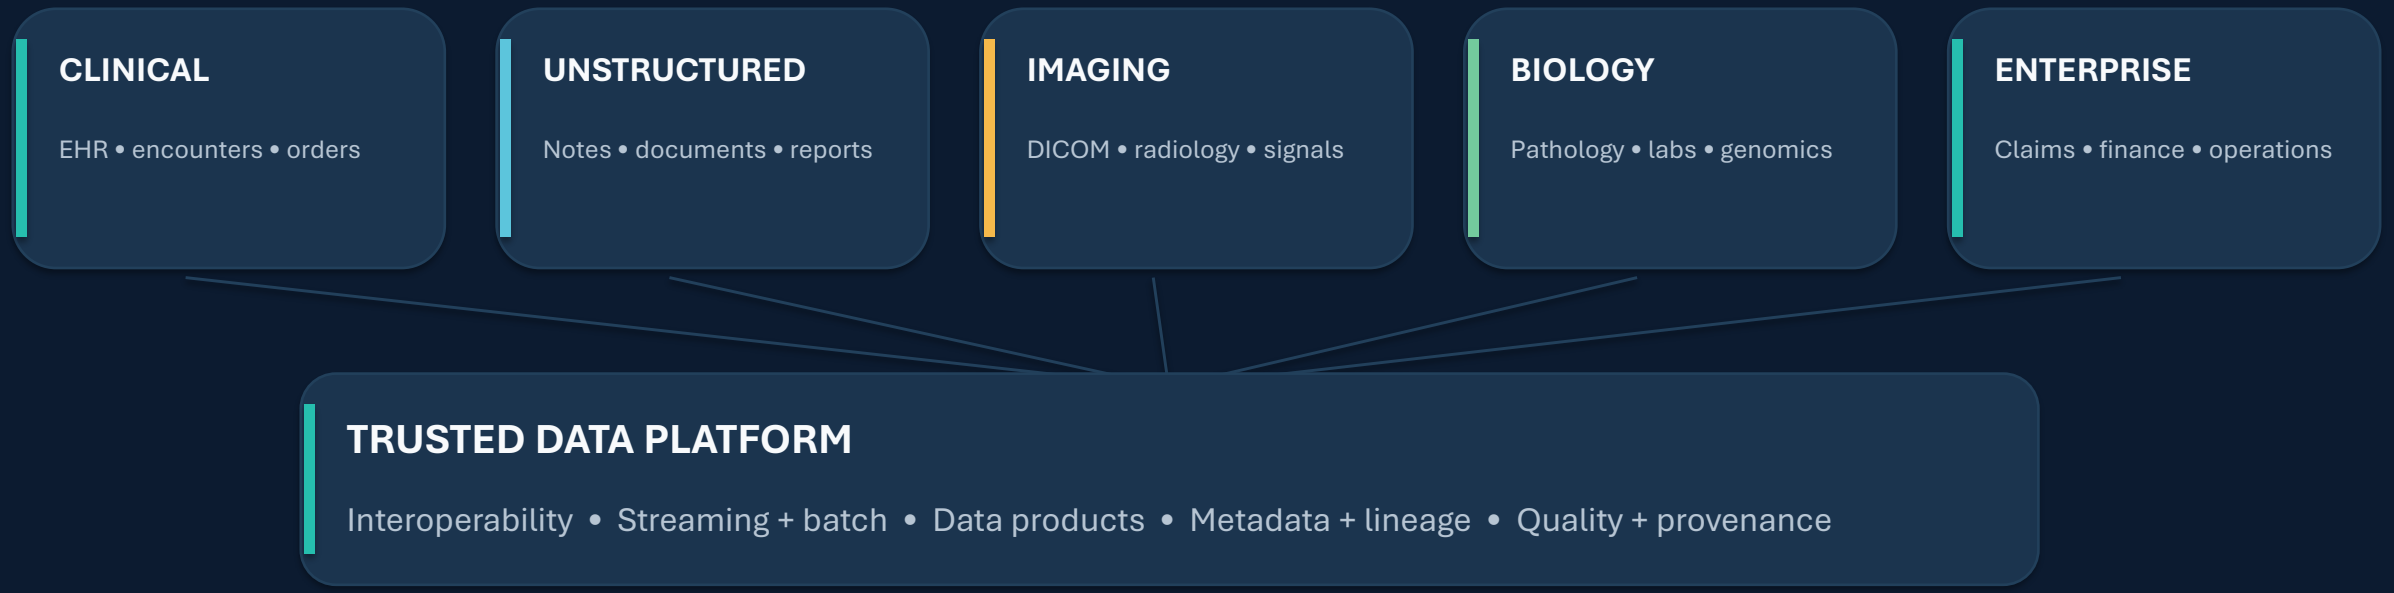
THE CONTEXT GAP

Siloed systems • Different vocabularies • Uneven data quality • Batch + real-time signals • Incomplete provenance

Integration makes information available. Architecture must also make it understandable and actionable.

Build a multimodal data foundation, not another silo

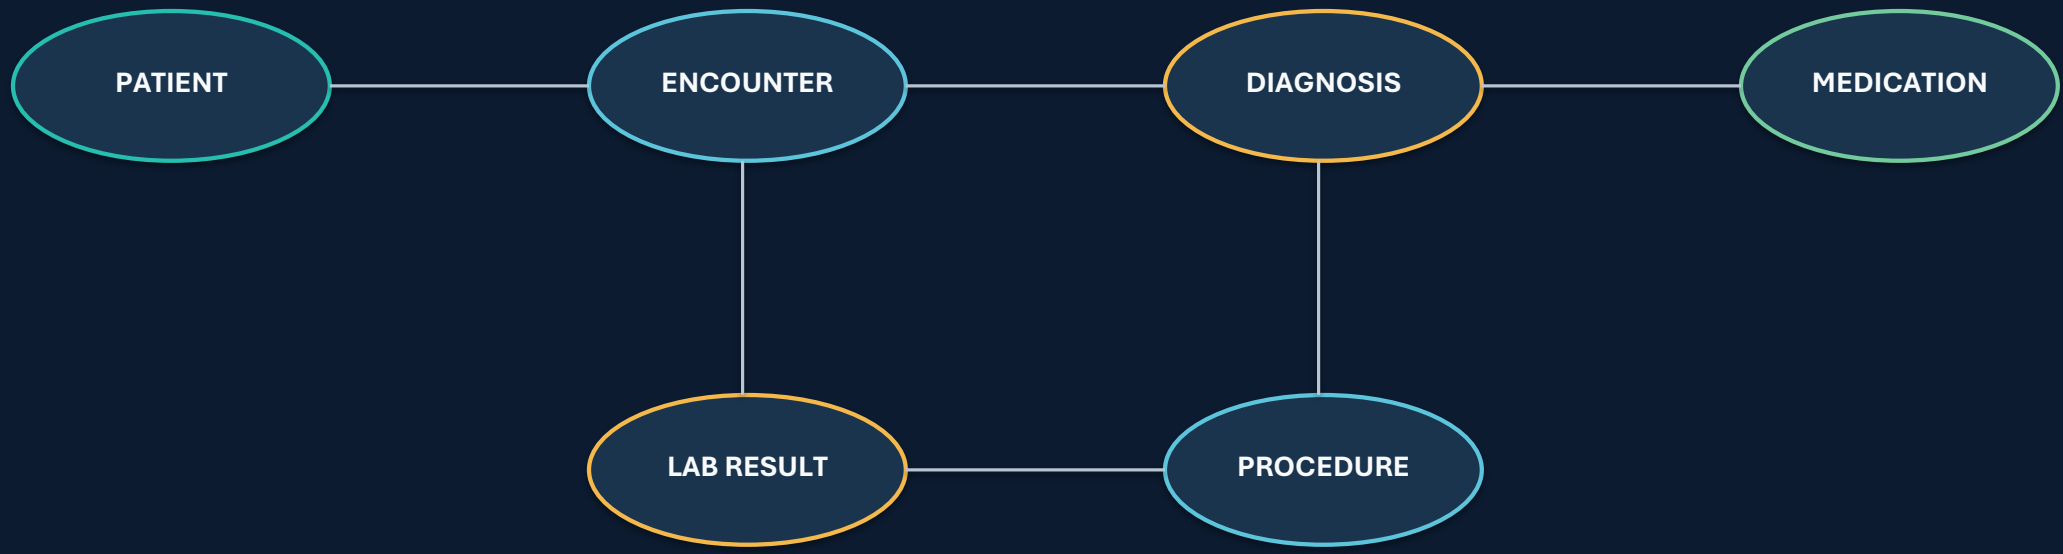
Make trusted information discoverable, connected, timely, and fit for purpose.



Design for access patterns and purpose: centralized, federated, and close-to-source where it makes sense.

The semantic layer gives agents healthcare context

Move from locating data to understanding what it means and how concepts relate.



Terminologies + ontologies | Knowledge graphs | Business rules | Domain definitions | Metadata

“Data integration makes information available. Semantic architecture makes information meaningful.”

Give agents a governed path from intent to execution

Context-aware agents combine retrieval, reasoning, tools, and workflow controls.

1 **UNDERSTAND**

User intent + task

2 **GROUND**

Retrieve relevant data + knowledge

3 **REASON**

Plan steps + assess constraints

4 **ORCHESTRATE**

Call approved APIs + tools

5 **VALIDATE**

Check results + escalate

TRADITIONAL ASSISTANT

Summarizes a patient record or answers a question.

CONTEXT-AWARE WORKFLOW

Prepares a review package and routes it for authorized approval.

An agent without enterprise context is an intelligent model operating with incomplete knowledge.

Govern the entire agent lifecycle, not just the model

Governance must travel with the agent from data access and reasoning through action and outcome.

IDENTITY + ACCESS

Agent identity • least privilege • scoped tools

DATA + KNOWLEDGE

Quality • lineage • provenance • purpose

MODEL + AGENT

Versioning • evaluation • drift • behavior

WORKFLOW + SAFETY

Human approval • escalation • safe boundaries

OBSERVABILITY

Logs • traces • tool calls • outcomes

“In healthcare, the question is not how much autonomy we can give an agent, but how much we can responsibly trust it with.”

A reusable architecture for governed Agentic AI

Build shared enterprise capabilities so use cases do not become disconnected point solutions.

EXPERIENCE + WORKFLOWS	Clinical • Research • Operations • Revenue cycle
AGENT ORCHESTRATION	Agent runtime • Planning • Memory/state • Human handoff • Multi-agent coordination
CONTEXT + KNOWLEDGE SERVICES	Semantic layer • Knowledge graph • RAG • Terminologies • Policy + rules
DATA + INTEROPERABILITY	EHR • Imaging • Pathology • Genomics • Claims • Events • APIs • Data products

GOVERNANCE

Cross-cutting: identity • security • privacy • evaluation • observability • audit • resilience

Start with bounded workflows and measurable value

The same foundations can support clinical, operational, and research intelligence.

ONCOLOGY

Connect clinical history, imaging, pathology, genomics, and treatment history to prepare a tumor-board review package.

HOSPITAL OPERATIONS

Combine census, ED status, bed capacity, staffing, and discharge signals to identify bottlenecks and coordinate recommendations.

RESEARCH

Support cohort discovery and study feasibility across clinical, genomic, and research data with privacy and access controls.

AUTONOMY SHOULD BE RISK-BASED

ASSIST

Summarize • retrieve

COORDINATE

Prepare • route

EXECUTE WITH APPROVAL

Human authorizes

BOUNDED AUTONOMY

Low-risk, predefined actions

Do not maximize autonomy everywhere. Match autonomy to clinical risk, impact, and demonstrated trust.

Build reusable foundations. Scale responsible intelligence.

Align clinical, data, AI, architecture, security, risk, and operations around outcomes.

CONNECTED DATA

Give agents access to the complete, trusted picture.

SEMANTIC CONTEXT

Help agents understand what information means.

EMBEDDED GOVERNANCE

Ensure agents act transparently within boundaries.

“We should not build AI agents around isolated use cases. We should build reusable enterprise capabilities that allow agents to operate across the healthcare ecosystem.”

The future will be defined not by how many agents we deploy, but by how intelligently, safely, and effectively they improve decisions, coordinate work, and support better outcomes.

THANK YOU • QUESTIONS